

An Improved Decision Procedure for Linear Time Mu-Calculus

Yao Liu, Zhenhua Duan and Cong Tian

Abstract—An improved Present Future form (PF form) for linear time μ -calculus (ν TL) is presented in this paper. In particular, the future part of the new version turns into the conjunction of elements in the closure of a formula. We show that every closed ν TL formula can be transformed into the new PF form. Additionally, based on the PF form, an algorithm for constructing Present Future form Graph (PFG), which can be utilized to describe models of a formula, is given. Further, an intuitive and efficient decision procedure for checking satisfiability of the guarded fragment of ν TL formulas based on PFG is proposed and implemented in C++. The new decision procedure has the best time complexity over the existing ones despite the cost of exponential space. Finally, a PFG-based model checking approach for ν TL is discussed where a counterexample can be obtained visually when a model violates a property.

Index Terms—Linear time μ -calculus, present future form graph, satisfiability, decision procedure, model checking.

1 INTRODUCTION

LINEAR time μ -calculus (ν TL) [1], linear time counterpart of modal μ -calculus [2], extends LTL [3] with least and greatest fixpoint operators. It is a formalism succinct in syntax and strong in expressive power which captures the expression of full ω -regular properties [4], [5]. Hence, it is useful for specifying and verifying various properties of concurrent programs and has received ever growing interest in the past few decades. From an application point of view, it is of great significance to establish a decision procedure for checking satisfiability of ν TL formulas. The work, however, is not easy due to the nesting of fixpoint operators.

Satisfiability and model checking [6] are two main decision problems for ν TL, which are both PSPACE-complete in complexity [7]. By satisfiability we denote the problem to find a decision procedure for determining whether a formula is satisfiable, while by model checking we mean the problem to decide whether all paths of a given Kripke structure satisfy a certain property. Moreover, decision procedures for checking satisfiability always play a critical role in deriving model checking approaches.

A lot of work has been done for achieving efficient decision procedures. The major milestone of the decision problems for modal μ -calculus is made by Streett and Emerson [8] who introduce the notion of well-founded pre-models and apply automata theory to check satisfiability. Related methods [9], [10] translate a formula into an equivalent alternating tree automaton and then check for emptiness. In [7], Vardi first adapts Streett and Emerson's method to ν TL with past operators which yields an algorithm running in $2^{O(|\phi|^4)}$. Later, Banieqbal and Barringer [11] show that if a formula has a model, then it is able to generate a good Hintikka structure which can be further transformed into a good path searching problem from a graph. Their algorithm is equivalent in time complexity to Vardi's but runs in exponential space. In [12], Stirling and Walker present a tableau characterisation for ν TL's decision problems without mentioning complexity issues. Bradfield, Esparza and Mader [13] improve the system of Stirling

and Walker based on the work in [14] by simplifying the success conditions for a tableau and their algorithm runs in $2^{O(|\phi|^2 \log |\phi|)}$. In [15], Dax, Hofmann and Lange propose a decision procedure for checking validity of ν TL formulas running in $2^{O(|\phi|^2 \log |\phi|)}$ and implemented in OCAML. To the best of our knowledge, all these existing decision methods mentioned above except for [15] are relatively complicated and concerned with theoretical aspects rather than practical applications. Therefore, we are motivated to formalize a more efficient and practical decision procedure.

To this end, a new Present Future form (PF form) for ν TL formulas is presented in this paper and we prove that every closed ν TL formula can be transformed into this form. Compared with [16], the definition of the new PF form, which still consists of the present and future parts, is more elegant: the present part remains unchanged while the future part turns into the conjunction of elements in the closure of a given formula. This further facilitates the proof of finiteness of Present Future form Graph (PFG) which can be used to describe models of a formula. A path in a PFG characterizes exactly a pre-model [7], [8] of the corresponding formula. Additionally, an algorithm, based on PF form, for constructing PFG is given. In a PFG, an edge may be associated with a mark which is a subset of variables occurring in the formula and utilized to keep track of the infinite unfolding problem for least fixpoint formulas. Further, a decision procedure for checking satisfiability of the guarded fragment of ν TL formulas based on PFG is presented. It is realized, with the help of marks, by searching for a ν -path in a PFG on which no least fixpoint formula unfolds itself infinitely. Moreover, the decision procedure has been implemented in C++. The result shows that our method improves the current best time complexity, $2^{O(|\phi|^2 \log |\phi|)}$ [13], [14], [15], to $2^{O(|\phi|)}$ despite the cost of exponential space.

According to the proposed decision procedure, a PFG-based model checking approach for ν TL is also proposed. To do so, first, an algorithm for constructing the product of a Kripke structure and a PFG is presented. Subsequently, we apply the notion of ν -paths in PFGs to the product graphs. Further, given a Kripke structure M and a desired property ϕ , the model checking approach is achieved by searching for a ν -path in the product graph of M and the PFG of $\neg\phi$. If such a path can be found, we will obtain a counterexample;

• The authors are with the ICTT and ISN Laboratory, Xidian University, Xi'an 710071, China. E-mail: yao_liu@stu.xidian.edu.cn, {zhhdian, ctian}@mail.xidian.edu.cn.

otherwise, M satisfies ϕ .

The idea of this paper is inspired by the normal form and normal form graph of Propositional Projection Temporal Logic (PPTL) [17], [18] which have played a vital role in obtaining a decision procedure for checking the satisfiability [19], [20], [21]. Compared with the existing methods for checking satisfiability of ν TL formulas, our decision procedure has the following advantages: (1) it does not depend on automata theory by considering PFGs; (2) it is more efficient in time and practical meanwhile; (3) it gives good insight into why and how a given formula is satisfiable through its PFG; (4) when a Kripke structure violates a property, it intuitively reflects that why a path is a counterexample through the corresponding product graph.

To summarize, our contributions are as follows:

- We define a new PF form for ν TL formulas and prove that every closed ν TL formula can be transformed into this form.
- We provide an algorithm for constructing PFG which can be used to describe models of a formula. During the constructing process, marks are technically added, which are useful in keeping track of the infinite unfolding problem for least fixpoint formulas.
- We introduce the notion of ν -paths and present a decision procedure for checking satisfiability of the guarded fragment of ν TL formulas by finding a ν -path in a PFG.
- We show that our decision procedure has the current best time complexity. We implement the decision procedure in C++ and experimental results show that our algorithm performs better than the one given in [15].
- We apply the notion of ν -paths in PFGs to the product graphs and propose a PFG-based model checking approach for ν TL.

The rest of this paper is organized as follows. The syntax and semantics of ν TL and some basic notions are introduced in Section 2. The new PF form of ν TL formulas is presented in Section 3. Section 4 describes an algorithm for constructing PFG and the decision procedure for checking satisfiability of the guarded fragment of ν TL formulas based on PFG is given in Section 5. Section 6 presents a model checking approach for ν TL based on PFG. Related work is discussed in section 7. Conclusions are drawn in Section 8.

2 PRELIMINARIES

2.1 Syntax and Semantics of ν TL

Let \mathcal{P} be a set of atomic propositions, and \mathcal{V} a set of variables. ν TL formulas are constructed based on the following syntax:

$$\phi ::= p \mid \neg p \mid X \mid \phi \vee \phi \mid \phi \wedge \phi \mid \bigcirc \phi \mid \mu X. \phi \mid \nu X. \phi$$

where p ranges over \mathcal{P} and X over \mathcal{V} .

We use σ to denote either μ or ν . An occurrence of a variable X in a formula is called *free* when it does not lie within the scope of σX ; it is called *bound* otherwise. A formula is called *closed* when it contains no free variables. Given two ν TL formulas ϕ_1 and ϕ_2 , we say $\phi_1 \leq \phi_2$ iff ϕ_2 is a subformula of ϕ_1 , and $\phi_1 < \phi_2$ iff ϕ_2 is a proper subformula of ϕ_1 . We write $\phi[\phi'/Y]$ for the result of simultaneously substituting ϕ' for all free occurrences of variable Y in ϕ . For each variable X in a formula, we assume that X is bound at most once. Thus, it can be seen that all formulas constructed by the syntax above are in positive normal form [22],

i.e. negations can be applied only to atomic propositions and each variable occurring in a formula is bound at most once.

For each bound variable X in formula ϕ , the unique subformula of ϕ in the form of $\sigma X. \varphi$ is said to be *identified by* X . The bound variables in ϕ can be partially ordered based on the nesting of their identified fixpoint formulas. Specifically, given two bound variables X and Y in ϕ , we say X is *higher than* Y iff the fixpoint formula identified by Y is a proper subformula of the one identified by X .

A formula is called a *guarded* one if, for each bound variable X in that formula, every occurrence of X is in the scope of a \bigcirc operator. Every formula can be transformed into an equivalent one in guarded form [23]. Note that the transformation causes an exponential increase in the size of a formula in the worst case [24].

Example 1. Translating formula $\nu X.(p \wedge \mu Y.(q \vee X \wedge \bigcirc Y))$ into guarded form.

$$\begin{aligned} & \nu X.(p \wedge \mu Y.(q \vee X \wedge \bigcirc Y)) \\ \equiv & \nu X.(p \wedge (q \vee X \wedge \bigcirc \mu Y.(q \vee X \wedge \bigcirc Y))) \\ & \text{by law } \sigma X. \phi \equiv \phi[\sigma X. \phi / X] \\ \equiv & \nu X.(p \wedge q \vee p \wedge \bigcirc \mu Y.(q \vee X \wedge \bigcirc Y)) \\ & \text{by law } \nu X.(X \wedge \phi \vee \varphi) \equiv \nu X.(\phi \vee \varphi) \end{aligned}$$

ν TL formulas are interpreted over linear time structures. A *linear time structure* over \mathcal{P} is a function $\mathcal{K}: \mathbb{N} \rightarrow 2^{\mathcal{P}}$ where \mathbb{N} denotes the set of natural numbers. The semantics of a ν TL formula ϕ , relative to \mathcal{K} and an environment $e: \mathcal{V} \rightarrow 2^{\mathbb{N}}$, is inductively defined as follows:

$$\begin{aligned} \llbracket p \rrbracket_e^{\mathcal{K}} &:= \{i \in \mathbb{N} \mid p \in \mathcal{K}(i)\} \\ \llbracket \neg p \rrbracket_e^{\mathcal{K}} &:= \{i \in \mathbb{N} \mid p \notin \mathcal{K}(i)\} \\ \llbracket X \rrbracket_e^{\mathcal{K}} &:= e(X) \\ \llbracket \varphi \vee \psi \rrbracket_e^{\mathcal{K}} &:= \llbracket \varphi \rrbracket_e^{\mathcal{K}} \cup \llbracket \psi \rrbracket_e^{\mathcal{K}} \\ \llbracket \varphi \wedge \psi \rrbracket_e^{\mathcal{K}} &:= \llbracket \varphi \rrbracket_e^{\mathcal{K}} \cap \llbracket \psi \rrbracket_e^{\mathcal{K}} \\ \llbracket \bigcirc \varphi \rrbracket_e^{\mathcal{K}} &:= \{i \in \mathbb{N} \mid i+1 \in \llbracket \varphi \rrbracket_e^{\mathcal{K}}\} \\ \llbracket \mu X. \varphi \rrbracket_e^{\mathcal{K}} &:= \bigcap \{W \subseteq \mathbb{N} \mid \llbracket \varphi \rrbracket_{e[X \mapsto W]}^{\mathcal{K}} \subseteq W\} \\ \llbracket \nu X. \varphi \rrbracket_e^{\mathcal{K}} &:= \bigcup \{W \subseteq \mathbb{N} \mid W \subseteq \llbracket \varphi \rrbracket_{e[X \mapsto W]}^{\mathcal{K}}\} \end{aligned}$$

where $e[X \mapsto W]$ is the environment e' agreeing with e except for $e'(X) = W$. e is used to evaluate free variables and can be dropped when ϕ is closed.

For a given formula ϕ , we say ϕ is true at state i of linear time structure \mathcal{K} , denoted by $\mathcal{K}, i \models \phi$, iff $i \in \llbracket \phi \rrbracket_e^{\mathcal{K}}$. We say ϕ is *valid*, denoted by $\models \phi$, iff $\mathcal{K}, j \models \phi$ for all linear time structures \mathcal{K} and all states j of \mathcal{K} ; ϕ is *satisfiable* iff there exists a linear time structure \mathcal{K} and a state j of \mathcal{K} such that $\mathcal{K}, j \models \phi$.

2.2 Approximant

Let *Ord* denote the class of ordinals. *Approximants* of fixpoint formulas are defined inductively by: $\mu^0 X. \phi = \perp$, $\nu^0 X. \phi = \top$, $\sigma^{\alpha+1} X. \phi = \phi[\sigma^\alpha X. \phi / X]$, $\mu^{\lambda} X. \phi = \bigvee_{\alpha < \lambda} \mu^\alpha X. \phi$ and $\nu^{\lambda} X. \phi = \bigwedge_{\alpha < \lambda} \nu^\alpha X. \phi$ where $\alpha, \lambda \in \text{Ord}$. In particular, λ is a limit ordinal.

The following lemma [25] is a standard result about approximants.

Lemma 1. For a linear time structure \mathcal{K} , we say $\mathcal{K}, 0 \models \nu X. \phi$ iff $\forall \alpha \in \text{Ord}$, $\mathcal{K}, 0 \models \nu^\alpha X. \phi$ and $\mathcal{K}, 0 \models \mu Y. \phi$ iff $\exists \alpha \in \text{Ord}$, $\mathcal{K}, 0 \models \mu^\alpha Y. \phi$. ([25])

Note that in both cases above, α is not a limit ordinal.

Let ϕ be a closed ν TL formula with exactly n μ -variables: X_1, \dots, X_n such that X_i is higher than X_j implies $i < j$. A μ -signature for ϕ is a tuple $\zeta = (\alpha_1, \dots, \alpha_n) \in (\mathbb{N} \cup \{\omega\})^n$ where each

α_i is an ordinal. A μ -signature *with respect to* a variable Y in ϕ is the prefix $(\alpha_1, \dots, \alpha_i)$ of ζ such that $Y = X_i$ when Y is a μ -variable, or X_i is the last μ -variable higher than Y when Y is a ν -variable. We write $\zeta(i)$ for the i -th component of ζ . For two μ -signatures ζ_1 and ζ_2 for ϕ , we write $\zeta_1 < \zeta_2$ to mean that ζ_1 lexicographically precedes ζ_2 , i.e. $\zeta_1(j) < \zeta_2(j)$ and $\zeta_1(i) = \zeta_2(i)$ for some j and each $i < j$. Note that the lexicographic ordering on μ -signatures is well-founded.

For a linear time structure \mathcal{K} and a state j of \mathcal{K} , we say $\mathcal{K}, j \models_\zeta \phi$ if (\mathcal{K}, j) is a model of ϕ resulting from ϕ with every least fixpoint subformula $\mu X_i. \phi_i$ of ϕ being interpreted by $\mu^{\zeta(i)} X_i. \phi_i$.

2.3 Closure

The *closure* $CL(\phi)$ of a formula ϕ , based on [26], is the least set of formulas such that

- (i) $\phi, \text{true} \in CL(\phi)$,
- (ii) if $\varphi \vee \psi$ or $\varphi \wedge \psi \in CL(\phi)$, then $\varphi \in CL(\phi)$ and $\psi \in CL(\phi)$,
- (iii) if $\bigcirc \varphi \in CL(\phi)$, then $\varphi \in CL(\phi)$,
- (iv) if $\sigma X. \varphi \in CL(\phi)$, then $\varphi[\sigma X. \varphi / X] \in CL(\phi)$.

Example 2. The closure of formula $\nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X)$.

$$\begin{aligned} CL(\nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X)) = \\ \{ \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X), \text{true}, \\ \mu Y. (\bigcirc Y \vee p \wedge \bigcirc \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X)), \\ \bigcirc \mu Y. (\bigcirc Y \vee p \wedge \bigcirc \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X)), \\ \vee p \wedge \bigcirc \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X), \\ \bigcirc \mu Y. (\bigcirc Y \vee p \wedge \bigcirc \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X)), \\ p \wedge \bigcirc \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X), \\ p, \bigcirc \nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X) \} \end{aligned}$$

It has been proved that the size of $CL(\phi)$ is linear in the size of ϕ (denoted by $|\phi|$) [26].

2.4 Dependency Relationship

Definition 1. For two formulas $\sigma X. \phi$ and $\sigma Y. \phi'$ where $\sigma X. \phi < \sigma Y. \phi'$, we say Y depends on X , denoted by $X \triangleleft Y$, iff X occurs free in ϕ' .

Note that the dependency relationship is transitive in a formula.

Example 3. Dependency relationship between variables.

- I. $\nu X. (\bigcirc X \wedge \mu Y. (p \vee \bigcirc Y))$
- II. $\nu X. \mu Y. (\bigcirc Y \vee p \wedge \bigcirc X) \vee \mu Z. \nu W. (\bigcirc Z \vee q \wedge \bigcirc W)$
- III. $\mu X. \nu Y. (\bigcirc X \vee \mu Z. \bigcirc (Z \vee Y \wedge p))$

In formula I, X and Y do not depend on each other. In formula II, we have $X \triangleleft Y$ and $Z \triangleleft W$, while in formula III we have $X \triangleleft Y \triangleleft Z$.

3 PF FORM OF ν TL FORMULAS

In this section, we first define PF form of ν TL formulas and then prove that every closed ν TL formula can be transformed into this form.

3.1 PF Form

Definition 2. Let ϕ be a closed ν TL formula, \mathcal{P}_ϕ the set of atomic propositions appearing in ϕ . PF form of ϕ is defined by:

$$\phi \equiv \bigvee_{i=1}^n (\phi_{p_i} \wedge \bigcirc \phi_{f_i})$$

where $\phi_{p_i} \equiv \bigwedge_{h=1}^{n_1} p_{ih}$, $p_{ih} \in \mathcal{P}_\phi$ for each h (i denotes either r or $\neg r$ for each $r \in \mathcal{P}_\phi$) and $\phi_{f_i} \equiv \bigwedge_{m=1}^{n_2} \phi_{im}$, $\phi_{im} \in CL(\phi)$ for each m .

The main difference between the PF form presented here and the one in [16] lies in the future part: in this paper, the future part is the conjunction of elements in the closure of a given formula rather than a closed formula in [16]. Thus, it can be seen that the PF form presented here is more rigorous in structure and this will dramatically simplify the proof of finiteness of PFG.

In the following, we prove that every closed ν TL formula can be transformed into PF form. For technical reasons, from now on we confine ourselves only to guarded formulas with no \vee appearing as the main operator under each \bigcirc operator. This can be easily achieved by pushing \bigcirc operators inwards using the equivalence $\bigcirc(\phi_1 \vee \phi_2) \equiv \bigcirc \phi_1 \vee \bigcirc \phi_2$.

Theorem 2. Every closed ν TL formula φ can be transformed into PF form.

Proof. Let $Conj(\psi)$ represent the set of all conjuncts in formula ψ . The proof proceeds by induction on the structure of φ .

- **Base case:**

- φ is p (or $\neg p$): p (or $\neg p$) can be transformed as:

$$p \equiv p \wedge \bigcirc \text{true} \text{ (or } \neg p \equiv \neg p \wedge \bigcirc \text{true})$$

The theorem holds obviously in these two cases.

- **Induction:**

- φ is $\bigcirc \phi$: $\bigcirc \phi$ can be written as:

$$\bigcirc \phi \equiv \bigvee_i (\text{true} \wedge \bigcirc \phi_i)$$

For each $\phi_c \in Conj(\phi_i)$, we have $\phi_c \in CL(\varphi)$ since $\phi \in CL(\varphi)$. Hence, φ can be transformed into PF form in this case.

- φ is $\phi_1 \vee \phi_2$: by induction hypothesis, both ϕ_1 and ϕ_2 can be transformed into PF form:

$$\phi_1 \equiv \bigvee_{i=1}^n (\phi_{1p_i} \wedge \bigcirc \phi_{1f_i}), \quad \phi_2 \equiv \bigvee_{j=1}^m (\phi_{2p_j} \wedge \bigcirc \phi_{2f_j})$$

where $\phi_{1c} \in Conj(\phi_{1f_i})$ and $\phi_{1c} \in CL(\phi_1)$, $\phi_{2c} \in Conj(\phi_{2f_j})$ and $\phi_{2c} \in CL(\phi_2)$, for each i and j . Then, we have

$$\varphi \equiv \phi_1 \vee \phi_2 \equiv \bigvee_{i=1}^n (\phi_{1p_i} \wedge \bigcirc \phi_{1f_i}) \vee \bigvee_{j=1}^m (\phi_{2p_j} \wedge \bigcirc \phi_{2f_j})$$

Since $\phi_1 \vee \phi_2 \in CL(\varphi)$, we have $\phi_1, \phi_2 \in CL(\varphi)$. For each $\phi_{1c} \in Conj(\phi_{1f_i})$, by induction hypothesis, we have $\phi_{1c} \in CL(\phi_1)$. Therefore, $\phi_{1c} \in CL(\varphi)$. Similarly, we can obtain that each $\phi_{2c} \in CL(\varphi)$. Thus, φ can be transformed into PF form in this case.

- φ is $\phi_1 \wedge \phi_2$: by induction hypothesis, both ϕ_1 and ϕ_2 can be transformed into PF form:

$$\phi_1 \equiv \bigvee_{i=1}^n (\phi_{1p_i} \wedge \bigcirc \phi_{1f_i}), \quad \phi_2 \equiv \bigvee_{j=1}^m (\phi_{2p_j} \wedge \bigcirc \phi_{2f_j})$$

where $\phi_{1c} \in Conj(\phi_{1f_i})$ and $\phi_{1c} \in CL(\phi_1)$, $\phi_{2c} \in Conj(\phi_{2f_j})$ and $\phi_{2c} \in CL(\phi_2)$, for each i and j . Then φ can be further converted into:

$$\begin{aligned} \varphi \equiv \phi_1 \wedge \phi_2 &\equiv \left(\bigvee_{i=1}^n (\phi_{1p_i} \wedge \bigcirc \phi_{1f_i}) \right) \wedge \left(\bigvee_{j=1}^m (\phi_{2p_j} \wedge \bigcirc \phi_{2f_j}) \right) \\ &\equiv \bigvee_{i=1}^n \bigvee_{j=1}^m (\phi_{1p_i} \wedge \phi_{2p_j} \wedge \bigcirc (\phi_{1f_i} \wedge \phi_{2f_j})) \end{aligned}$$

Since $\phi_1 \wedge \phi_2 \in CL(\varphi)$, we have $\phi_1, \phi_2 \in CL(\varphi)$. For each $\phi_{1c} \in Conj(\phi_{1f_i})$, by induction hypothesis, we have $\phi_{1c} \in CL(\phi_1)$. Hence, $\phi_{1c} \in CL(\varphi)$. Similarly, we can obtain that each $\phi_{2c} \in CL(\varphi)$. Therefore, all conjuncts behind \bigcirc operators in φ belong to $CL(\varphi)$ and φ can be transformed into PF form in this case.

- φ is $\mu X.\phi$: let p_X be an atomic proposition where $\llbracket p_X \rrbracket^K = \llbracket \mu X.\phi \rrbracket^K$ w.r.t. a certain linear time structure \mathcal{K} . As a result, $\phi[p_X/X]$ can be treated as a closed formula. By induction hypothesis, $\phi[p_X/X]$ can be transformed into PF form:

$$\phi[p_X/X] \equiv \bigvee_{i=1}^n (\phi_{p_i} \wedge \bigcirc \phi_{f_i}[p_X/X])$$

Due to the restriction of guarded form, p_X can only appear in the future part of the above PF form. Suppose

$$U_1 = \{\phi_{f_1}, \dots, \phi_{f_m}\}, U_2 = \{\phi_{f_{m+1}}, \dots, \phi_{f_n}\}$$

where each $\phi_{f_j} \in U_1$ ($j \in \{1, \dots, m\}$) does not contain p_X while each $\phi_{f_k} \in U_2$ ($k \in \{m+1, \dots, n\}$) contains p_X . By induction hypothesis, for each $\phi_{c_j} \in Conj(\phi_{f_j})$ and $\phi_{c_k}[p_X/X] \in Conj(\phi_{f_k})$, we have $\phi_{c_j}, \phi_{c_k}[p_X/X] \in CL(\phi[p_X/X])$. Since $\mu X.\phi \equiv \phi[\mu X.\phi/X]$, then φ can be converted into:

$$\varphi \equiv \phi[\mu X.\phi/X] \equiv \bigvee_{i=1}^n (\phi_{p_i} \wedge \bigcirc \phi_{f_i}[\mu X.\phi/p_X])$$

For each $\phi_{c_j} \in CL(\phi[p_X/X])$, after the substitution of $\mu X.\phi$ for p_X , we can still have $\phi_{c_j} \in CL(\phi[\mu X.\phi/p_X])$. Since $\phi[\mu X.\phi/p_X] \equiv \phi[\mu X.\phi/X]$ and $\phi[\mu X.\phi/X] \in CL(\varphi)$, then $\phi_{c_j} \in CL(\varphi)$. For each $\phi_{c_k}[p_X/X] \in CL(\phi[p_X/X])$, after the substitution of $\mu X.\phi$ for p_X , we can further obtain $\phi_{c_k}[\mu X.\phi/p_X] \in CL(\phi[\mu X.\phi/p_X])$. Since $\phi[\mu X.\phi/p_X] \equiv \phi[\mu X.\phi/X]$ and $\phi[\mu X.\phi/X] \in CL(\varphi)$, we have $\phi_{c_k}[\mu X.\phi/p_X] \in CL(\varphi)$. Therefore, φ can be transformed into PF form in this case.

- φ is $\nu X.\phi$: this case can be proved similarly to the case φ is $\mu X.\phi$.

Thus, it can be concluded that every closed ν TL formula can be transformed into PF form. \square

3.2 Algorithm for Transforming a ν TL Formula Into PF Form

In this section we present algorithm *PFTran* for transforming a closed ν TL formula ϕ into PF form. The basic idea of the algorithm comes directly from the proof of Theorem 2. Thus, its correctness can be ensured.

Algorithm 1 *PFTran*(ϕ)

- 1: **case**
 - 2: ϕ is *true*: **return** *true* $\wedge \bigcirc$ *true*
 - 3: ϕ is *false*: **return** *false*
 - 4: ϕ is ϕ_p where $\phi_p \equiv \bigwedge_{h=1}^n \dot{p}_h$: **return** $\phi_p \wedge \bigcirc$ *true*
 - 5: ϕ is $\phi_p \wedge \bigcirc \varphi$: **return** $\bigvee_i (\phi_p \wedge \bigcirc \varphi_i)$
 - 6: ϕ is $\bigcirc \varphi$: **return** $\bigvee_i (\text{true} \wedge \bigcirc \varphi_i)$
 - 7: ϕ is $\phi_1 \vee \phi_2$: **return** $PFTran(\phi_1) \vee PFTran(\phi_2)$
 - 8: ϕ is $\phi_1 \wedge \phi_2$: **return** $AND(PFTran(\phi_1), PFTran(\phi_2))$
 - 9: ϕ is $\sigma X.\varphi$: **return** $PFTran(\varphi[\sigma X.\varphi/X])$
 - 10: **end case**
-

In algorithm *PFTran*, if ϕ is *true* or *false*, the transformation is straightforward; if ϕ is ϕ_p where $\phi_p \equiv \bigwedge_{h=1}^n \dot{p}_h$, its PF form is $\phi_p \wedge \bigcirc$ *true*; if ϕ is $\phi_p \wedge \bigcirc \varphi$, its PF form is $\bigvee_i (\phi_p \wedge \bigcirc \varphi_i)$; if ϕ is $\bigcirc \varphi$, its PF form is $\bigvee_i (\text{true} \wedge \bigcirc \varphi_i)$; if ϕ is $\phi_1 \vee \phi_2$, the algorithm calls itself to transform ϕ_1 and ϕ_2 into PF form respectively; if ϕ is $\phi_1 \wedge \phi_2$, the algorithm also calls itself first to transform ϕ_1 and ϕ_2 into PF form respectively and then converts $\phi_1 \wedge \phi_2$ into PF form by algorithm *AND*; if ϕ is $\sigma X.\varphi$, the algorithm transforms $\varphi[\sigma X.\varphi/X]$ into PF form.

Algorithm 2 *AND*(ϕ, φ)

- 1: **if** ϕ is of the form $\bigvee_i (\phi_i \wedge \bigcirc \phi'_i)$ and φ is of the form $\bigvee_j (\varphi_j \wedge \bigcirc \varphi'_j)$ **then**
 - 2: **return** $\bigvee_i \bigvee_j (\phi_i \wedge \varphi_j \wedge \bigcirc (\phi'_i \wedge \varphi'_j))$
 - 3: **end if**
-

Algorithm *AND* is used by *PFTran* to deal with the \wedge construct. Note that the inputs ϕ and φ for *AND* are both in PF form. Therefore, ϕ must be of the form $\bigvee_i (\phi_i \wedge \bigcirc \phi'_i)$ while φ of the form $\bigvee_j (\varphi_j \wedge \bigcirc \varphi'_j)$.

In the following, we use an example to demonstrate how to transform a closed ν TL formula into PF form by means of algorithm *PFTran*.

Example 4. Transforming formula $\nu X.(r \wedge \bigcirc X) \wedge \mu Y.(q \vee p \wedge \bigcirc Y)$ into PF form by algorithm *PFTran*.

$$\begin{aligned}
& PFTran(\nu X.(r \wedge \bigcirc X) \wedge \mu Y.(q \vee p \wedge \bigcirc Y)) \\
& \equiv AND(PFTran(\nu X.(r \wedge \bigcirc X)), PFTran(\mu Y.(q \vee p \wedge \bigcirc Y))) \\
& \equiv AND(PFTran(r \wedge \bigcirc \nu X.(r \wedge \bigcirc X)), PFTran(q \vee p \wedge \bigcirc \mu Y.(q \vee p \wedge \bigcirc Y))) \\
& \equiv AND(r \wedge \bigcirc \nu X.(r \wedge \bigcirc X), PFTran(q) \vee PFTran(p \wedge \bigcirc \mu Y.(q \vee p \wedge \bigcirc Y))) \\
& \equiv AND(r \wedge \bigcirc \nu X.(r \wedge \bigcirc X), q \wedge \bigcirc \text{true} \vee p \wedge \bigcirc \mu Y.(q \vee p \wedge \bigcirc Y)) \\
& \equiv r \wedge q \wedge \bigcirc \nu X.(r \wedge \bigcirc X) \vee r \wedge p \wedge \bigcirc (\nu X.(r \wedge \bigcirc X) \wedge \mu Y.(q \vee p \wedge \bigcirc Y))
\end{aligned}$$

PF form enables us to convert a formula ϕ into two parts: the present and future ones. The present part is a conjunction of atomic propositions or their negations in ϕ , while the future part is a next formula consisting of the conjunction of formulas in $CL(\phi)$. To make ϕ satisfied, the present part should be satisfied at the current state while the future part at the next one. Further, we can repeat the transformation process by converting each formula in the future part into PF form, which inspires us to construct a graph, namely Present Future form Graph (PFG), for describing models of ϕ . This will be discussed in the next section.

4 PRESENT FUTURE FORM GRAPH

4.1 Definition of PFG

For a closed ν TL formula ϕ , the PFG of ϕ , denoted by G_ϕ , is a tuple (N_ϕ, E_ϕ, n_0) where N_ϕ is a set of nodes, E_ϕ a set of directed edges, and n_0 the root node. Each node in N_ϕ is specified by the conjunction of formulas in $CL(\phi)$ while each edge in E_ϕ is identified by a triple (ϕ_0, ϕ_e, ϕ_1) , where $\phi_0, \phi_1 \in N_\phi$ and ϕ_e is the label of the edge from ϕ_0 to ϕ_1 . An edge may be associated with a mark which is a subset of variables occurring in ϕ .

Definition 3. For a given closed ν TL formula ϕ , N_ϕ and E_ϕ can be inductively defined by:




Fig. 1. An example of PFG

- 1) $n_0 = \phi \in N_\phi$;
- 2) For all $\varphi \in N_\phi \setminus \{false\}$, if $\varphi \equiv \bigvee_{i=1}^k (\varphi_{p_i} \wedge \bigcirc \varphi_{f_i})$, then $\varphi_{f_i} \in N_\phi$, $(\varphi, \varphi_{p_i}, \varphi_{f_i}) \in E_\phi$ for each i ($1 \leq i \leq k$).

In a PFG, the root node is denoted by a double circle while each of other nodes by a single circle. Each edge is denoted by a directed arc with a label and also possibly a mark that connects two nodes. To simplify matters, we usually use variables to represent the corresponding fixpoint formulas occurring in a node. An example of PFG for formula $\mu X.(p \vee \bigcirc X) \vee \nu Y.(q \wedge \bigcirc Y)$ is depicted in Fig. 1. There are four nodes in the PFG where n_0 is the root node. (n_0, q, n_3) is an edge with label being q and mark being $\{Y\}$ while (n_0, p, n_1) is an edge with label being p and no mark.

4.2 Marks in PFG

From Fig. 1 we can see that there may exist a path in a PFG, e.g. $n_0, true, (n_2, true)^\omega$, which arises from the infinite unfolding of a least fixpoint formula. Thus, marks are useful in a PFG to keep track of the infinite unfolding problem for least fixpoint formulas when constructing the PFG.

Definition 4. Given a PFG G_ϕ and a node $\phi_m \in N_\phi$ where $\phi_m \equiv \bigvee_{i=1}^k (\phi_{p_i} \wedge \bigcirc \phi_{f_i})$. The mark of edge $(\phi_m, \phi_{p_i}, \phi_{f_i})$ ($1 \leq i \leq k$) is a set of variables M_v such that for each $X \in M_v$, the fixpoint formula $\sigma X.\phi_X$ identified by X appears as a subformula of ϕ_{f_i} and has not been unfolded by formula $\nu Y.\phi_Y$ where Y is higher than X in the PF form transformation process.

We use the notion of μ -signatures to demonstrate how to add marks to a PFG. Intuitively, a variable X is added to a mark in a PF form transformation process if the unfolding of the corresponding formulas does not increase the μ -signature w.r.t. X . As a result, we can use marks to detect the infinite descending chains of μ -signatures.

When transforming a formula into its PF form, the occurrence of a fixpoint formula $\sigma X.\phi_X$ in the future part ϕ_{f_i} may be caused by the unfolding of: (I) itself, (II) a least fixpoint formula $\mu Y.\phi_Y$ where Y is higher than X , or (III) a greatest fixpoint formula $\nu Z.\phi_Z$



Fig. 2. PFG of $\nu Z. \bigcirc (\bigcirc (\mu Y.(p \vee \bigcirc Y)) \wedge Z)$ 

Fig. 3. How the marks work

where Z is higher than X . According to Lemma 3.5 in [8], the μ -signature w.r.t. X does not increase unless the case III happens. For example, as shown in Fig. 2, when node n_0 is transformed into PF form: $n_0 \equiv true \wedge \bigcirc n_1$, the occurrence of $\mu Y.(p \vee \bigcirc Y)$ in n_1 is due to the unfolding of $\nu Z. \bigcirc (\bigcirc (\mu Y.(p \vee \bigcirc Y)) \wedge Z)$, hence Y does not exist in the mark of edge $(n_0, true, n_1)$.

Note that cases I and II, or I and III (e.g. the occurrence of Y in the mark of edge $(n_2, true, n_2)$ in Fig. 2) can occur simultaneously. If that happens, we can see that the μ -signature w.r.t. X still does not increase. In particular, cases II and III cannot happen simultaneously since X is bound exactly once.

Given a formula $\sigma X.\phi$, to construct its PFG sometimes we need to deal with a formula of the form $\bigwedge_{i=1}^n \sigma_i X_i.\phi_i$, where each $\sigma_i X_i.\phi_i \in CL(\sigma X.\phi)$ and $i < j$ implies X_i is higher than X_j , in a PF form transformation process. It is straightforward that the unfolding of $\sigma_n X_n.\phi_n$ ensures that the μ -signature w.r.t. each X_i does not increase after the transformation despite the value of each σ_i .

In the following, we use a simple example to illustrate how the marks work.

Example 5. Tracing the infinite unfolding of $\mu X.(p \vee \bigcirc X)$ using marks.

The PF form of $\mu X.(p \vee \bigcirc X)$ is: $\mu X.(p \vee \bigcirc X) \equiv p \wedge \bigcirc true \vee true \wedge \bigcirc \mu X.(p \vee \bigcirc X)$. The second disjunct of the PF form leads to the generation of edge $(n_0, true, n_0)$ in Fig. 3. Then mark $\{X\}$ is added accordingly since $\mu X.(p \vee \bigcirc X)$ appears in the future part of the PF form and has been unfolded by itself in the PF form transformation process. Moreover, it is easy to see that all other edges have no marks. Formula $\mu X.(p \vee \bigcirc X)$ indicates that the atomic proposition p finally holds somewhere and therefore path $(n_0, true)^\omega$ does not characterize a model. Actually, $(n_0, true)^\omega$ is generated by the infinite unfolding of $\mu X.(p \vee \bigcirc X)$ and the infinite occurrence of mark $\{X\}$ on this path describes exactly an infinite descending chain of μ -signatures for $\mu X.(p \vee \bigcirc X)$. This is why we need to use marks.

4.3 Paths in PFG

A path Π in a PFG G_ϕ is an infinite alternate sequence of nodes and edges departing from the root node. In the following, we show how to establish the relationship between paths in PFG and linear time structures.

Let $Atom(\bigwedge_{i=1}^m \dot{q}_i)$ denote the set of atomic propositions or their negations appearing in formula $\bigwedge_{i=1}^m \dot{q}_i$. Given a path $\Pi = \phi_0, \phi_{e0}, \phi_1, \phi_{e1}, \dots$ in a PFG, we can obtain a corresponding linear time structure $Atom(\phi_{e0}), Atom(\phi_{e1}), \dots$

Example 6. Paths in Fig. 4.

- 1) Path $n_0, true, (n_1, p)^\omega$ corresponds to the linear time structure $\{true\}\{p\}^\omega$.
- 2) Path $n_0, true, (n_2, true, n_1, p)^\omega$ corresponds to the linear time structure $\{true\}(\{true\}\{p\})^\omega$.




Fig. 4. Paths in PFG

Actually, each node in the PFG G_ϕ of formula ϕ corresponds precisely to a consistent subset of $CL(\phi)$. In other words, each path in G_ϕ characterizes a pre-model [7], [8] of ϕ . A pre-model is almost a model except that it ignores the infinite unfolding problem for least fixpoint formulas. We can distinguish real models from all pre-models using marks.

4.4 Algorithm for Constructing PFG

Given a closed ν TTL formula ϕ , the whole process of constructing its PFG G_ϕ is presented in Algorithm 3.

Algorithm 3 PFGCon(ϕ)

```

1:  $n_0 = \phi$ ,  $N_\phi = \{n_0\}$ ,  $E_\phi = \emptyset$ ,  $isHandled[n_0] = 0$ 
2: while there exists  $\varphi \in N_\phi \setminus \{false\}$  and  $isHandled[\varphi] = 0$  do
3:    $isHandled[\varphi] = 1$ 
4:    $\varphi = PFTran(\varphi)$  /*suppose  $\varphi = \bigvee_{i=1}^k (\varphi_{p_i} \wedge \bigcirc \varphi_{f_i})$ */
5:   for  $i = 1$  to  $k$  do
6:     if  $\varphi_{p_i}$  is not false then
7:        $E_\phi = E_\phi \cup \{(\varphi, \varphi_{p_i}, \varphi_{f_i})\}$  /*adding edges*/
8:        $AddMark((\varphi, \varphi_{p_i}, \varphi_{f_i}))$  /*obtaining the corresponding marks*/
9:     if  $\varphi_{f_i} \notin N_\phi$  then
10:       $N_\phi = N_\phi \cup \{\varphi_{f_i}\}$  /*adding nodes*/
11:      if  $\varphi_{f_i}$  is not false then
12:         $isHandled[\varphi_{f_i}] = 0$  /* $\varphi_{f_i}$  is a new node which needs to be handled*/
13:      else
14:         $isHandled[\varphi_{f_i}] = 1$  /* $\varphi_{f_i}$  does not need to be handled*/
15:      end if
16:    end if
17:  end for
18: end while
19: for all  $\varphi \in N_\phi$  with no outgoing edge do
20:    $N_\phi = N_\phi \setminus \{\varphi\}$  /*eliminating redundant nodes and the relative edges*/
21:    $E_\phi = E_\phi \cup \{(\varphi_i, \varphi_e, \varphi)\}$ 
22: end for
23: return  $G_\phi$ 

```

The algorithm takes ϕ as input and returns G_ϕ . First, n_0 is assigned to ϕ . N_ϕ and E_ϕ are initialized to $\{n_0\}$ and empty, respectively. Further, the algorithm repeatedly converts an unhandled formula $\varphi \in N_\phi$ into PF form by algorithm $PFTran$ and then adds the corresponding nodes and edges to N_ϕ and E_ϕ , respectively, until all formulas in N_ϕ have been handled. $isHandled[]$ is used to indicate whether a formula has been handled. If $isHandled[\varphi] = 0$, φ needs further to be handled; otherwise, φ has been handled or there is no need to handle it. Function $AddMark$ is utilized to mark

an edge with a subset of variables occurring in ϕ by distinguishing appropriate fixpoint formulas from all fixpoint formulas contained in the future part of a certain PF form.

Algorithm 4 AddMark($(\varphi, \varphi_{p_i}, \varphi_{f_i})$)

```

1: for each conjunct  $\varphi_c$  of  $\varphi_{f_i}$  do
2:   if  $\varphi_c$  is of the form  $\bigcirc^n \sigma X. \varphi_X$  and  $flag[\sigma X. \varphi_X] = 0$  then
3:      $M_i = M_i \cup \{X\}$  /* $M_i$  represents the mark of edge  $(\varphi, \varphi_{p_i}, \varphi_{f_i})$ */
4:   end if
5: end for

```

The input for algorithm $AddMark$ is an edge $(\varphi, \varphi_{p_i}, \varphi_{f_i})$ in G_ϕ . In the algorithm, $flag[]$ is employed to denote whether a fixpoint formula φ_{fix} appearing in the future part of a PF form has been unfolded by a greatest fixpoint formula $\nu Y. \varphi_Y$ where Y is higher than the variable identifying φ_{fix} in the PF form transformation process. If $flag[\varphi_{fix}] = 1$, φ_{fix} is unfolded by $\nu Y. \varphi_Y$; otherwise, it is unfolded by itself or a least fixpoint formula. For any fixpoint subformula $\sigma Z. \varphi_{sub}$ of φ , $flag[\sigma Z. \varphi_{sub}]$ is assigned to 0 before φ is transformed into PF form. For the input $(\varphi, \varphi_{p_i}, \varphi_{f_i})$, $AddMark$ checks each conjunct φ_c of φ_{f_i} . If φ_c is in the form $\bigcirc^n \sigma X. \varphi_X$ ($n \geq 0$) and $flag[\sigma X. \varphi_X] = 0$, X is added to M_i . Here \bigcirc^n represents the consecutive occurrence of \bigcirc operators for n times and M_i represents the mark of the edge $(\varphi, \varphi_{p_i}, \varphi_{f_i})$.

Additionally, it is worth pointing out that, throughout the construction of G_ϕ , a *false* node (e.g. $p \wedge \neg p$) may be generated which corresponds to an inconsistent subset of $CL(\phi)$. Such kind of nodes have no successor and are redundant. We use the *for loop* in Line 20 of $PFGCon$ to remove those redundant nodes as well as the relative edges.

Example 7. Constructing the PFG of formula $\mu X. (p \vee \bigcirc X) \vee \nu Y. (q \wedge \bigcirc Y)$ by algorithm $PFGCon$.

As depicted in Fig. 1, at the very beginning, the root node $\mu X. (p \vee \bigcirc X) \vee \nu Y. (q \wedge \bigcirc Y)$ is created and denoted by n_0 ; then we transform $\mu X. (p \vee \bigcirc X) \vee \nu Y. (q \wedge \bigcirc Y)$ into PF form:

$$\mu X. (p \vee \bigcirc X) \vee \nu Y. (q \wedge \bigcirc Y) \equiv p \wedge \bigcirc true \vee true \wedge \bigcirc \mu X. (p \vee \bigcirc X) \vee q \wedge \bigcirc \nu Y. (q \wedge \bigcirc Y)$$

Accordingly, nodes *true*, $\mu X. (p \vee \bigcirc X)$ and $\nu Y. (q \wedge \bigcirc Y)$ are created and denoted respectively by n_1 , n_2 and n_3 . Meanwhile, edges (n_0, p, n_1) , $(n_0, true, n_2)$ and (n_0, q, n_3) are created among which $(n_0, true, n_2)$ is marked with $\{X\}$ and (n_0, q, n_3) with $\{Y\}$. Further, *true* is transformed into PF form: $true \equiv true \wedge \bigcirc true$. Thus, edge $(n_1, true, n_1)$ is created. After that, $\mu X. (p \vee \bigcirc X)$ is transformed into PF form: $\mu X. (p \vee \bigcirc X) \equiv p \wedge \bigcirc true \vee true \wedge \bigcirc \mu X. (p \vee \bigcirc X)$. Hence, edges (n_2, p, n_1) and $(n_2, true, n_2)$ are created where $(n_2, true, n_2)$ is marked with $\{X\}$. Finally, $\nu Y. (q \wedge \bigcirc Y)$ is transformed into PF form: $\nu Y. (q \wedge \bigcirc Y) \equiv q \wedge \bigcirc \nu Y. (q \wedge \bigcirc Y)$. Accordingly, edge (n_3, q, n_3) is created with the mark being $\{Y\}$, and the whole construction process terminates.

4.5 Finiteness of PFG

In the PFG G_ϕ of formula ϕ generated by algorithm $PFGCon$, N_ϕ and E_ϕ are produced by repeatedly transforming the unhandled nodes into PF form. Since each node in N_ϕ is the conjunction of formulas in $CL(\phi)$, the following corollary is easily obtained.

Corollary 3. For any closed ν TL formula ϕ , the number of nodes in G_ϕ is bounded by $2^{O(|\phi|)}$, and the number of edges in G_ϕ is bounded by $2^{O(|\phi|)} \cdot 2^{O(n_p)} \cdot 2^{O(n_v)} \cdot 2^{O(|\phi|)}$ (which is also $2^{O(|\phi|)}$), where n_p and n_v denote the number of atomic propositions and fixpoint subformulas occurring in ϕ respectively.

5 DECISION PROCEDURE BASED ON PFG

In this section we show how to find a model for a given closed ν TL formula ϕ from its PFG G_ϕ . In fact, each outgoing edge of a node in G_ϕ amounts to a possible choice prescribed by an underlying choice function [8]. Since a node cannot have multiple choices simultaneously, we restrict ourselves here only to paths ending with simple loops in G_ϕ . Let Π be a path in G_ϕ , for convenience, we use $LES(\Pi)$ to denote the set of edges appearing in the loop part of Π , $Mark(E)$ the mark of edge E , $LMS(\Pi)$ the set of all μ -variables occurring in each $Mark(E_i)$ where $E_i \in LES(\Pi)$.

5.1 ν -path

Here we present the notion of ν -paths which will play a vital role in obtaining the PFG-based decision procedure for ν TL.

Definition 5. Given a PFG G_ϕ and a path Π in G_ϕ , we call Π a ν -path iff for each $X \in LMS(\Pi)$, an edge $E \in LES(\Pi)$ can be found such that $X \notin Mark(E)$ and there exists no $X' \in Mark(E)$ where $X \triangleleft X'$.

Example 8. ν -paths in Fig. 5.

- 1) $\Pi_1: (n_0, p \wedge q)^\omega$. Π_1 is a ν -path since $LMS(\Pi_1) = \emptyset$.
- 2) $\Pi_2: n_0, true, (n_1, true)^\omega$. We have $LES(\Pi_2) = \{(n_1, true, n_1)\}$ and $LMS(\Pi_2) = \{Y, W\}$. For the first variable $Y \in LMS(\Pi_2)$, we cannot find an edge from $LES(\Pi_2)$ whose mark does not contain Y . So Π_2 is not a ν -path.
- 3) $\Pi_3: n_0, p, (n_2, true, n_1, p)^\omega$. We have $LES(\Pi_3) = \{(n_2, true, n_1), (n_1, p, n_2)\}$ and $LMS(\Pi_3) = \{Y, W\}$. For the first variable $Y \in LMS(\Pi_3)$, we can find an edge $(n_1, p, n_2) \in LES(\Pi_3)$ whose mark does not contain Y and any variable depending on Y . However, for the second variable $W \in LMS(\Pi_3)$, we cannot find an edge from $LES(\Pi_3)$ whose mark does not contain W . Therefore, Π_3 is not a ν -path.
- 4) $\Pi_4: (n_0, q, n_3, p \wedge q)^\omega$. We have $LES(\Pi_4) = \{(n_0, q, n_3), (n_3, p \wedge q, n_0)\}$ and $LMS(\Pi_4) = \{Y\}$. For the only variable $Y \in LMS(\Pi_4)$, we can find an edge $(n_3, p \wedge q, n_0) \in LES(\Pi_4)$




Fig. 5. ν -paths in PFG

whose mark does not contain Y and any variable depending on Y . Thus, Π_4 is a ν -path.

- 5) $\Pi_5: n_0, p, (n_2, true, n_1, q, n_3, true, n_1, p)^\omega$. We have $LES(\Pi_5) = \{(n_2, true, n_1), (n_1, q, n_3), (n_3, true, n_1), (n_1, p, n_2)\}$ and $LMS(\Pi_5) = \{Y, W\}$. For the first variable $Y \in LMS(\Pi_5)$, we can find an edge $(n_1, p, n_2) \in LES(\Pi_5)$ whose mark does not contain Y and any variable depending on Y . Further, for the second variable $W \in LMS(\Pi_5)$, we can find an edge $(n_1, q, n_3) \in LES(\Pi_5)$ whose mark does not contain W and any variable depending on W . Therefore, Π_5 is a ν -path.

Regarding the notion of ν -paths, the following theorem is formalized.

Theorem 4. A closed ν TL formula ϕ is satisfiable iff a ν -path can be found in G_ϕ .

Proof. (\Rightarrow) Suppose ϕ is satisfiable and no ν -path exists in G_ϕ . In this case, for any path Π_1 in G_ϕ , there exists at least one $X \in LMS(\Pi_1)$ such that for each edge $E_1 \in LES(\Pi_1)$, either $X \in Mark(E_1)$ or $X' \in Mark(E_1)$, where $X \triangleleft X'$.

As a result, we can obtain the following sequence of variables according to the sequence of marks in the loop part of Π_1 :

$$X, X_1, X_2, \dots, X_n, X$$

where each X_i ($1 \leq i \leq n$) is either X itself or a variable depending on X .

Further, according to the sequence of variables above, we can acquire the following sequence of fixpoint formulas:

$$\mu X.\phi_X, \sigma X_1.\phi_1, \sigma X_2.\phi_2, \dots, \sigma X_n.\phi_n, \mu X.\phi_X$$

Here each $\sigma X_i.\phi_i$ is identified by X_i and $\mu X.\phi_X$ by X . Since each X_i is either X or a variable depending on X , $\mu X.\phi_X$ must appear as a subformula of each $\sigma X_i.\phi_i$. According to the way the marks are added, it can be seen that the above sequence describes exactly an infinite descending chain of μ -signatures w.r.t. X . By the well-foundedness of μ -signatures we can derive that Π_1 does not characterize a model of ϕ . This contradicts the premise that ϕ is satisfiable. Therefore, if ϕ is satisfiable, there exists at least one ν -path in G_ϕ .

(\Leftarrow) Let Π_2 be a ν -path in G_ϕ .

When $LMS(\Pi_2)$ is empty, no infinite descending chain of μ -signatures on Π_2 can be detected. Consequently, Π_2 characterizes a model of ϕ .

When $LMS(\Pi_2)$ is not empty, we have that for each $Y \in LMS(\Pi_2)$, an edge $E_2 \in LES(\Pi_2)$ can be found such that $Y \notin Mark(E_2)$ and there exists no $Y' \in Mark(E_2)$ where $Y \triangleleft Y'$. Subsequently, for each sequence of variables relevant to Y obtained according to the sequence of marks in the loop part of Π_2 :

$$Y, Y_1, Y_2, \dots, Y_m, Y$$

We can obtain the following sequence of fixpoint formulas:

$$\mu Y.\phi_Y, \sigma Y_1.\phi_1, \sigma Y_2.\phi_2, \dots, \sigma Y_m.\phi_m, \mu Y.\phi_Y$$

where there must exist a formula $\sigma Y_j.\phi_j$ ($1 \leq j \leq m$) in which $\mu Y.\phi_Y$ does not appear as a subformula. Similarly, we have that Π_2 characterizes a model of ϕ according to the well-foundedness of μ -signatures w.r.t. Y . It follows that when there exists a ν -path in G_ϕ , ϕ is satisfiable. \square

Consequently, we reduce the satisfiability problem of ν TL formulas to a ν -path searching problem from a PFG.




Fig. 6. Examples of PFGs for satisfiability checking

Example 9. Checking satisfiability of the following formulas.

- (1) $\nu Z.(\mu X.(\bigcirc X \vee \nu Y.(p \wedge \bigcirc Y)) \wedge \bigcirc Z)$
- (2) $\nu X.(p \wedge \bigcirc X) \wedge \nu Y.(\neg p \wedge \bigcirc Y)$
- (3) $\mu X.(\mu Y.(p \wedge \bigcirc Y) \vee \bigcirc X)$

For formula (1), as depicted in Fig. 6 (a), since a ν -path $n_0, p, (n_2, p, n_3, p)^\omega$ can be found in its PFG, it is satisfiable. For formula (2), as its PFG is empty and contains no ν -path, it is unsatisfiable. For formula (3), as depicted in Fig. 6 (b), no ν -path exists in its PFG and hence it is unsatisfiable.

5.2 Implementation of the Decision Procedure

Based on Theorem 4, a PFG-based decision procedure, algorithm *PFGSAT*, for checking satisfiability of ν TL formulas is derived.

Algorithm 5 PFGSAT(ϕ)

```

1:  $G_\phi = \text{PFGCon}(\phi)$ 
2: if  $G_\phi$  is empty then
3:   return unsatisfiable
4: end if
5:  $\text{Tarjan}(G_\phi, n_0)$ 
6: for each  $\text{scc} \in \text{sccs}$  do
7:    $\text{SCCNuSearch}(v, \text{scc})$  /*  $v$  is an arbitrary node in  $\text{scc}$  */
8: end for
9: return unsatisfiable

```

The algorithm takes a closed ν TL formula ϕ as input and returns the result whether ϕ is satisfiable in the end. To do so, the PFG, G_ϕ , of ϕ is constructed first. Next, it checks whether G_ϕ is empty: if so, ϕ is unsatisfiable since no ν -path can be found in G_ϕ ; otherwise, the algorithm will try to find a ν -path in G_ϕ . Further, algorithm *Tarjan* is employed to compute all strongly connected components (SCCs) in G_ϕ . Finally, the algorithm checks whether there exists a loop in some SCC which corresponds to a ν -path by algorithm *SCCNuSearch*: if so, *SCCNuSearch* will return that ϕ is satisfiable; otherwise, ϕ is unsatisfiable.

SCC Computation. Tarjan algorithm [27] presented in Algorithm 6 is a classical algorithm for computing SCCs in a graph based on depth-first search (DFS). The algorithm takes a PFG G_ϕ and a node v in G_ϕ as inputs and acquires all SCCs in G_ϕ . $\text{dfn}[u]$ is employed to represent the timestamp of a given node u indicating the number of nodes which have been visited before u is visited, while $\text{low}[u]$ the timestamp of the earliest node reachable from u or subtrees of u . Also, we use $\text{visit}[u]$ to denote whether a node u has been visited. If $\text{visit}[u] = 1$, u has already been visited; otherwise, u has not been visited yet. For each node u in G_ϕ , $\text{visit}[u]$ is initialized to 0. $\text{src}[]$ and $\text{tgt}[]$ are utilized to obtain the source and target nodes of an edge, respectively.

Algorithm 6 Tarjan(G_ϕ, v)

```

1:  $\text{dfn}[v] = \text{low}[v] = ++\text{index}$ 
2:  $\text{visit}[v] = 1$ 
3:  $\text{Stack.push}(v)$ 
4: for each edge  $e \in E_\phi$  where  $\text{src}[e] = v$  do
5:   if  $\text{visit}[\text{tgt}[e]] = 0$  then
6:      $\text{Tarjan}(G_\phi, \text{tgt}[e])$ 
7:      $\text{low}[v] = \min\{\text{low}[v], \text{low}[\text{tgt}[e]]\}$ 
8:   else
9:     if  $\text{tgt}[e]$  is in  $\text{Stack}$  then
10:       $\text{low}[v] = \min\{\text{low}[v], \text{dfn}[\text{tgt}[e]]\}$ 
11:    end if
12:   end if
13: end for
14: if  $\text{dfn}[v] = \text{low}[v]$  then
15:    $\text{subGraph scc}$ 
16:   repeat
17:      $u = \text{Stack.top}()$ 
18:      $\text{Stack.pop}()$ 
19:      $\text{scc.push\_back}(u)$ 
20:   until  $v = u$ 
21:    $\text{sccs.push\_back}(\text{scc})$ 
22: end if

```

Algorithm 7 SCCNuSearch(v, scc)

```

1:  $\text{NS.push\_back}(v)$ 
2: for each edge  $e$  in  $\text{scc}$  do
3:   if  $\text{src}[e] = v$  and  $\text{visit}[e] = 0$  then
4:      $\text{ES.push\_back}(e)$ 
5:      $\text{visit}[e] = 1$ 
6:     if  $\text{isLoop}(\text{tgt}[e], \text{pos})$  then
7:        $\text{TES.assign}(\text{ES.begin}() + \text{pos}, \text{ES.end}())$ 
8:       if  $\text{isNuPath}(\text{TES})$  then
9:         return satisfiable
10:      end if
11:       $\text{ES.pop\_back}()$ 
12:     else
13:        $\text{SCCNuSearch}(\text{tgt}[e], \text{scc})$ 
14:     end if
15:   end if
16: end for
17: if  $\text{ES.size}() > 0$  then
18:    $\text{ES.pop\_back}()$ 
19: end if
20:  $\text{NS.pop\_back}()$ 

```

Path Construction. Given an SCC scc in a PFG G_ϕ and an arbitrary node v in scc , we use algorithm *SCCNuSearch* to build a path which is likely to correspond to a ν -path in G_ϕ . Two global variables, ES and NS , are used in the algorithm. ES is a vector which stores the sequence of edges aiming to construct a path ending with a loop. NS is also a vector storing the sequence of nodes corresponding to ES . In addition, $\text{src}[]$ and $\text{tgt}[]$ are employed to obtain the source and target nodes of an edge, respectively. The algorithm uses $\text{visit}[]$ to indicate whether an edge e has been visited. If $\text{visit}[e] = 1$, e has already been visited; otherwise, e has not been visited yet. For each edge e in G_ϕ , $\text{visit}[e]$ is initialized to 0. isLoop and isNuPath are two boolean functions. isLoop determines whether a node u exists in




Fig. 8. PFG of formula II

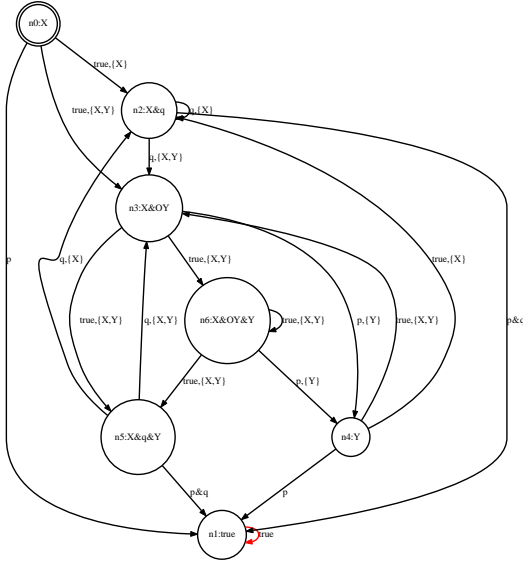


Fig. 9. PFG of formula III




Fig. 10. PFG of formula IV




Fig. 11. PFG of formula V

5.3 Experimental Results

We have implemented a prototype of our PFG-based decision procedure in C++. Given a guarded formula, the prototype is able to construct its PFG and find a ν -path from the PFG. To evaluate the performance of our tool, we compare it with the tool given in [15] which is the only tool available for the decision problems of ν TL.

In [15], the authors have checked the validity of the following three families of formulas: *Include_n*, *Nester_n* and *Counter_n* on a 1G memory PC.

$$\text{Include}_n \equiv \nu X. \underbrace{(q \wedge \bigcirc(q \wedge \bigcirc(\dots \bigcirc(q \wedge \bigcirc(\neg q \wedge \bigcirc X)) \dots)))}_{2n \text{ times}} \rightarrow \nu Z. \mu Y. (\neg q \wedge \bigcirc Z \vee q \wedge \bigcirc(q \wedge \bigcirc Y))$$

Nester_n $\equiv \psi \vee \neg\psi$ where

$$\psi \equiv \mu X_1. \nu X_2. \mu X_3. \dots \sigma X_n. (q_1 \vee \bigcirc(X_1 \wedge (q_2 \vee \bigcirc(X_2 \wedge \dots (q_n \vee \bigcirc X_n) \dots))))$$

$$\text{Counter}_n \equiv \bigvee_{i=0}^n \neg c_i \vee \mu X. (\bigcirc X \vee (c_0 \leftrightarrow \bigcirc \neg c_0) \vee \bigvee_{i=1}^n (\bigcirc c_i \leftrightarrow c_i \wedge \neg c_{i-1} \vee c_{i-1} \wedge (\bigcirc c_{i-1} \leftrightarrow c_i)))$$

Include_n describes the property $((aa)^nb)^\omega \subseteq ((aa)^*b)^\omega$, where the alphabet symbol a is the label $\{q\}$ and b is \emptyset . Note that *Include_n*

TABLE 1
Experimental results

n	$\neg Include_n$			$\neg Nester_n$			$\neg Counter_n$		
	Time (ms)	PFG nodes (number)	PFG edges (number)	Time (ms)	PFG nodes (number)	PFG edges (number)	Time (ms)	PFG nodes (number)	PFG edges (number)
0	0	6	18	—	—	—	0	2	2
1	31	17	39	0	1	1	0	4	4
2	63	28	64	31	10	30	16	8	8
3	124	39	85	1,185	73	386	156	16	16
4	218	50	106	128,559	601	4,640	889	32	32
5	328	61	127	18,075,924	5,401	55,419	5,117	64	64

is not LTL-definable for any $n \in \mathbb{N}$. $Nester_n$ is a class of formulas with several alternating fixpoint operators. $\neg Counter_n$ formalizes an $(n + 1)$ -bit counter.

We equivalently check satisfiability of $\neg Include_n$, $\neg Nester_n$ and $\neg Counter_n$. The experiments are carried out on a 1.73GHz, Genuine Intel(R) CPU T2080 with 1G of memory. Table 1 presents the empirical measures for complexity of the PFG-based decision procedure. The columns *Time* denote the running time to decide satisfiability of each formula. The columns *PFG nodes* (resp. *PFG edges*) represent the number of nodes (resp. edges) in the PFG of the corresponding formula. In [15], the running time for checking validity of each formula is always around a few minutes. In addition, they suffer from the problem of memory overflow for formulas $Nester_4$, $Nester_5$ and $Counter_5$. However, the satisfiability in most cases can be decided in less than 1 second using our tool. In particular, it takes only about 5 seconds to decide satisfiability of $\neg Counter_5$, while the satisfiability of $\neg Nester_4$ and $\neg Nester_5$ can be decided in about 2 minutes and 5 hours, respectively. Therefore, it can be seen that our method has a better performance in practice.

5.4 Complexity Issues

In this section we discuss the complexity of the PFG-based decision procedure. Let ϕ be a closed vTL formula, $G_\phi = (N_\phi, E_\phi, n_0)$ the PFG of ϕ , n_v the number of fixpoint subformulas appearing in ϕ . We write $|\phi|$ for the size of ϕ , $|N_\phi|$ and $|E_\phi|$ for the number of nodes and edges in G_ϕ , respectively. We can obtain, by Corollary 3, that both $|N_\phi|$ and $|E_\phi|$ are bounded by $2^{O(|\phi|)}$. Regarding ϕ , we have the following lemmas.

Lemma 5. *Algorithm $PFTran$ can be done in $2^{O(|\phi|)}$.*

Proof. First of all, it can be seen that the running time of $PFTran$ depends mainly on the number of recursive calls for itself as well as the running time of algorithm AND .

The proof proceeds by induction on the structure of ϕ .

- **Base case:**
 - $\phi = true, false, \phi_p, \phi_p \wedge \bigcirc \varphi$ (where ϕ_p is of the form $\bigwedge_{h=1}^n \dot{p}_h$), or $\bigcirc \varphi$: the lemma holds obviously in these cases.
- **Induction:**
 - $\phi = \phi_1 \vee \phi_2$: by induction hypothesis, $PFTran(\phi_1)$ and $PFTran(\phi_2)$ can be finished in $2^{O(|\phi_1|)}$ and $2^{O(|\phi_2|)}$, respectively. Further, we can see that the running time of $PFTran(\phi)$ is $2^{O(|\phi_1|)} + 2^{O(|\phi_2|)}$, which is bounded by $2^{O(|\phi|)}$.
 - $\phi = \phi_1 \wedge \phi_2$: by induction hypothesis, $PFTran(\phi_1)$ and $PFTran(\phi_2)$ can be completed in $2^{O(|\phi_1|)}$ and $2^{O(|\phi_2|)}$, respectively. After being transformed into PF form, the number of disjuncts in ϕ_1 (resp. ϕ_2) is bounded by $2^{O(|\phi_1|)}$ (resp. $2^{O(|\phi_2|)}$). Hence, algorithm AND can be completed in

$2^{O(|\phi_1|+|\phi_2|)}$. Further, we can obtain that the overall running time of $PFTran(\phi)$ is $2^{O(|\phi_1|)} + 2^{O(|\phi_2|)} + 2^{O(|\phi_1|+|\phi_2|)}$, which is bounded by $2^{O(|\phi|)}$.

- $\phi = \sigma X.\varphi$: we consider only guarded formulas where each free occurrence of X in φ must be in the scope of a \bigcirc operator. Regarding X as an atomic proposition, φ can be transformed into PF form by algorithm $PFTran$, which can be accomplished, by induction hypothesis, in $2^{O(|\varphi|)}$. Subsequently, by substituting $\sigma X.\varphi$ for all free occurrences of X in φ (which can be done in linear time), we can obtain that $\varphi[\sigma X.\varphi/X]$ can also be transformed into PF form by algorithm $PFTran$ in $2^{O(|\varphi|)}$. Therefore, the running time of $PFTran(\sigma X.\varphi)$ is bounded by $2^{O(|\phi|)}$ in this case.

It follows that algorithm $PFTran$ can be done in $2^{O(|\phi|)}$. \square

Lemma 6. *Algorithm $PFGCon$ can be done in $2^{O(|\phi|)}$.*

Proof. The running time of $PFGCon$ depends mainly on three parts: (I) generating nodes and edges; (II) adding marks; (III) eliminating redundant nodes and the relative edges. In part I, since $|N_\phi|$ is bounded by $2^{O(|\phi|)}$, the number of iterations in Line 2 is bounded by $2^{O(|\phi|)}$. In each iteration, algorithm $PFTran$ is called, which can be finished in $2^{O(|\phi|)}$ according to Lemma 5. Next, after the PF form transformation, we can see that the number of iterations in Line 5 of $PFGCon$ is bounded by $2^{O(|\phi|)}$. Hence, part I can be finished in $2^{O(|\phi|)}$. In part II, $|E_\phi|$ is bounded by $2^{O(|\phi|)}$. For each edge in E_ϕ , we need to obtain its mark information. Algorithm $AddMark$ checks if a fixpoint formula, which has been unfolded by itself or a least fixpoint formula in a PF form transformation process, exists in the future part of the PF form and can be completed in $O(|\phi|)$. Therefore, part II can be completed in $2^{O(|\phi|)}$. Further, part III can apparently be finished in $2^{O(|\phi|)}$. Thus, based on the above analysis, the overall running time of $PFGCon$ is in $2^{O(|\phi|)}$. \square

Lemma 7. *Algorithm $Tarjan$ can be done in $2^{O(|\phi|)}$. ([27])*

Lemma 8. *Algorithm $isNuPath$ can be done in $2^{O(|\phi|)}$.*

Proof. In Line 1, the number of iterations is bounded by $2^{O(|\phi|)}$. In Line 2, the conditional statements $X \in Mark(e)$ and X is a μ -variable can be determined in $O(n_v)$ and $O(1)$, respectively. Further, the number of iterations in Line 6 (resp. Line 7) is bounded by $O(n_v)$ (resp. $2^{O(|\phi|)}$). For each μ -variable V_μ appearing in ϕ , we maintain a list of variables depending on V_μ . In this way, the conditional statement in Line 8 can be decided in $O(n_v^2)$. Therefore, algorithm $isNuPath$ can be done in $2^{O(|\phi|)}$. \square

Lemma 9. *Algorithm $SCCNuSearch$ can be done in $2^{O(|\phi|)}$.*

Proof. Since each edge in the input scc is handled exactly once, the total number of recursive calls for $SCCNuSearch$ is bounded

by $2^{O(|\phi|)}$. In Line 2, the number of iterations is also bounded by $2^{O(|\phi|)}$. Subsequently, algorithm *isLoop* and *isNuPath* are called as the conditional statements in Lines 6 and 8 and both of them can be determined in $2^{O(|\phi|)}$. It follows that algorithm *SCCNuSearch* can be done in $2^{O(|\phi|)}$. \square

Theorem 10. *The decision procedure PFGSAT can be done in $2^{O(|\phi|)}$.*

Proof. This theorem is a direct consequence of Lemmas 5-9. \square

As far as we know, the current best time complexity for the decision problems of ν TL is $2^{O(|\phi|^2 \log |\phi|)}$ due to [13], [14], [15] and our PFG-based decision procedure noticeably improves it to $2^{O(|\phi|)}$. However, the price to pay for the improvement is exponential space.

Remarks. Building the PFG for a given formula is similar to the process of constructing the tableau for that formula. The main difference is that marks are technically added during the PFG construction. How to add marks is guided by the condition that whether or nor the unfolding of the corresponding formulas will increase the μ -signature with respect to some variable in a PF form transformation process. As a result, we can detect non-well-foundedness of the unfolding of least fixpoint formulas within a PFG. The existing decision procedures, i.e. [15], usually need to construct an automaton to check non-well-foundedness. Therefore, the complexity of those methods is mainly influenced by the results from automata theory. With our method, the satisfiability of a formula can be simply decided through the PFG of the formula. Since our method is independent of the results from automata theory, we obtain a faster decision procedure. However, we have to use the information of the whole PFG when deciding satisfiability of a formula, hence our method can no longer be done in polynomial space.

6 MODEL CHECKING BASED ON PFG

In this section, we use Kripke structures as models to demonstrate how the PFG-based model checking approach for ν TL is achieved.

6.1 Kripke Structure

Let AP be a set of atomic propositions. A *Kripke structure* [28] over AP is defined as a quadruple $M = (S, s_0, R, I)$ consisting of:

- a finite set of states S ,
- a designated initial state $s_0 \in S$,
- a transition relation $R \subseteq S \times S$ where R is *total*, i.e. $\forall s \in S, \exists s' \in S, (s, s') \in R$,
- a labeling (or interpretation) function $I : S \rightarrow 2^{AP}$ defining for each state $s \in S$ the set of all atomic propositions valid in s .

A *path* of M is an infinite sequence of states $\rho = s_0, s_1, s_2, \dots$, departing from the initial state s_0 , such that for each $i \geq 0$, $(s_i, s_{i+1}) \in R$. The word on ρ is the sequence of sets of atomic propositions $w = I(s_0), I(s_1), I(s_2), \dots$ which is an ω -word over alphabet 2^{AP} .

We need to take all paths in a Kripke structure into consideration in terms of the model checking problem for ν TL. Given a Kripke structure M and a property ϕ specified by a ν TL formula, we say $M \models \phi$ iff every path in M satisfies ϕ . However, when determining whether $M \models \phi$, for simplicity, we usually check whether there exists a path in M satisfying $\neg\phi$: if not so, we have $M \models \phi$; otherwise, $M \not\models \phi$ and we can obtain a counterexample.

In the previous section, we have presented a decision procedure for checking satisfiability of the guarded fragment of ν TL formulas based on PFG. Therefore, according to the decision procedure, we are able to formalize a PFG-based model checking approach for ν TL. To do so, first, it is essential to construct the product of a Kripke structure and a PFG.

6.2 Product Graph

Let $M = (S, s_0, R, I)$ be a Kripke structure, $G_\phi = (N_\phi, E_\phi, n_0)$ the PFG of formula ϕ , and AP the set of atomic propositions over M and ϕ . The *product* of M and G_ϕ is defined as a triple $G_{M \times \phi} = (V, E, v_0)$ where:

- $V \subseteq S \times N_\phi$ is a set of nodes.
- $E \subseteq V \times Q \times V$ is a set of edges, where Q is the label of an edge between two nodes. Each $((s_i, \varphi_i), Q_i, (s_j, \varphi_j)) \in E$ satisfies three conditions: (1) $(s_i, s_j) \in R$, $(\varphi_i, \varphi_e, \varphi_j) \in E_\phi$; (2) $Q_i \equiv \varphi_e$; (3) $((s_i, \varphi_i), Q_i, (s_j, \varphi_j))$ has the same mark as $(\varphi_i, \varphi_e, \varphi_j)$.
- In particular, $v_0 = (s_0, n_0)$ is the root node.

In a product graph $G_{M \times \phi}$, a node is called a *dead node* if it has no outgoing edge. A *finite path* $\Omega = V_0, Q_0, V_1, Q_1, \dots, V_k$ in $G_{M \times \phi}$ is a finite alternate sequence of nodes and edges starting from the root node while ending with a dead node. An *infinite path* $\Omega = V_0, Q_0, V_1, Q_1, \dots$ in $G_{M \times \phi}$ is an infinite alternate sequence of nodes and edges departing from the root node.

Given a Kripke structure $M = (S, s_0, R, I)$ and the PFG $G_\phi = (N_\phi, E_\phi, n_0)$ of a formula ϕ , we use algorithm *PGConstruction* to construct their product.

Algorithm 10 PGConstruction(M, G_ϕ)

```

1:  $v_0 = (s_0, n_0)$ ,  $V = \{v_0\}$ ,  $E = \emptyset$ ,  $h[v_0] = 0$ 
2: while there exists  $v = (s, \varphi) \in V$  and  $h[v] = 0$  do
3:    $h[v] = 1$ 
4:   for each  $(s, s') \in R$  and  $(\varphi, \varphi_e, \varphi') \in E_\phi$  do
5:     if LabelCheck( $s, \varphi_e$ ) then
6:        $E = E \cup \{(s, \varphi), \varphi_e, (s', \varphi')\}$  /*the newly added edge
       has the same mark as edge  $(\varphi, \varphi_e, \varphi')$  in  $G_\phi$ */
7:       if  $(s', \varphi') \notin V$  then
8:          $V = V \cup \{(s', \varphi')\}$ 
9:          $h[(s', \varphi')] = 0$ 
10:      end if
11:    end if
12:  end for
13: end while
14: return  $G_{M \times \phi}$ 

```

The algorithm takes M and G_ϕ as inputs and returns the product graph $G_{M \times \phi}$ in the end. The root node, v_0 , of $G_{M \times \phi}$ is assigned to (s_0, n_0) . Moreover, the set of nodes V and the set of edges E in $G_{M \times \phi}$ are initialized to $\{v_0\}$ and empty, respectively. The algorithm repeatedly checks whether the construction could proceed on an unhandled node $v \in V$ using boolean function *LabelCheck*, and then adds, if so, the corresponding nodes and edges to V and E , respectively, until all nodes in V have been handled. $h[]$ is utilized to indicate whether a node has been handled. If $h[v] = 0$, v needs to be further handled; otherwise, v has already been handled.

Algorithm 11 LabelCheck(s, φ_e)

```

1: if  $\varphi_e$  is true then
2:   return true
3: end if
4: for each conjunct  $\dot{p}$  of  $\varphi_e$  do
5:   if  $\dot{p}$  is  $p$  and  $p \in I(s)$ , or  $\dot{p}$  is  $\neg p$  and  $p \notin I(s)$  then
6:     continue
7:   else
8:     return false
9:   end if
10: end for
11: return true

```

Given a node $v = (s, \varphi) \in V$, we use algorithm *LabelCheck* to determine if the construction could continue from a transition $(s, s') \in R$ and an edge $(\varphi, \varphi_e, \varphi') \in E_\phi$. If $\varphi_e \equiv \text{true}$, the output of *LabelCheck* is true; otherwise, for each conjunct \dot{p} of φ_e , where \dot{p} denotes an atomic proposition or its negation, if \dot{p} is p and $p \in I(s)$ (resp. \dot{p} is $\neg p$ and $p \notin I(s)$), the output of *LabelCheck* is true. In all other cases, the output of *LabelCheck* is false. Note that when an edge $((s, \varphi), \varphi_e, (s', \varphi'))$ is added to E , it retains the mark of edge $(\varphi, \varphi_e, \varphi')$ in G_ϕ .

Similar to the representation of a PFG, in a product graph, we also use a double circle to denote the root node and a single circle to denote each of other nodes. Each edge is denoted by a directed arc connecting two nodes. A mark is placed behind the label of an edge if it exists.



Example 11. Constructing the product of Kripke structure M_0 and the PFG of formula $\phi_0: \bigcirc \nu X.(p \wedge \bigcirc X)$ by algorithm *PGConstruction*.

As illustrated in Fig. 13, at the very beginning, the root node (s_0, n_0) is created and handled first by the algorithm. For the transition (s_0, s_1) in M_0 and the edge (n_0, true, n_1) in G_{ϕ_0} , since the label of (n_0, true, n_1) is *true*, the output of algorithm *LabelCheck* is *true*. Therefore, node (s_1, n_1) and edge $((s_0, n_0), \text{true}, (s_1, n_1))$ are created. Similarly, for the transition (s_0, s_3) and the edge (n_0, true, n_1) , node (s_3, n_1) and edge $((s_0, n_0), \text{true}, (s_3, n_1))$ are created.

Next, the algorithm deals with the node (s_1, n_1) . For the transition (s_1, s_2) in M_0 and the edge (n_1, p, n_1) in G_{ϕ_0} , since $p \in I(s_1)$, the output of algorithm *LabelCheck* is *true*. Therefore, node (s_2, n_1) and edge $((s_1, n_1), p, (s_2, n_1))$ are created. Moreover, $((s_1, n_1), p, (s_2, n_1))$ is marked with $\{X\}$.

Subsequently, the node (s_3, n_1) is handled by the algorithm. For the transition (s_3, s_3) in M_0 and the edge (n_1, p, n_1) in G_{ϕ_0} , since $p \in I(s_3)$, the output of algorithm *LabelCheck* is *true*. Thus, edge $((s_3, n_1), p, (s_3, n_1))$ is created and marked with $\{X\}$.

Further, the algorithm deals with the node (s_2, n_1) . For the transition (s_2, s_1) in M_0 and the edge (n_1, p, n_1) in G_{ϕ_0} , since $p \notin I(s_2)$, the output of algorithm *LabelCheck* is *false*, which indicates that the construction cannot proceed on node (s_2, n_1) . By now,

Fig. 12. Kripke structure M_0 and the PFG of ϕ_0 Fig. 13. Product of M_0 and G_{ϕ_0}

all nodes have been handled and the whole construction process terminates.

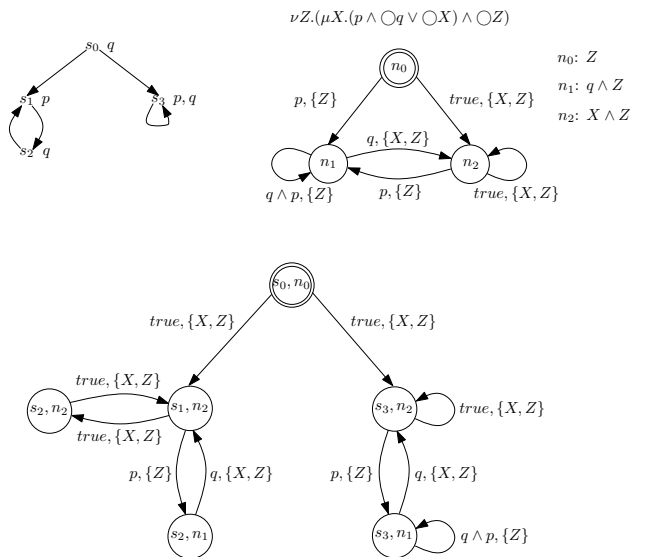
6.3 ν -paths in Product Graph

To formalize the PFG-based model checking approach for ν TTL, we apply the definition of ν -paths in PFGs to the product graphs. Similarly, we concentrate only on paths ending with loops in a product graph. Given an infinite path Ω in a product graph, for convenience, we use $LES_{MC}(\Omega)$ to denote the set of edges appearing in the loop part of Ω , $Mark_{MC}(e)$ the mark of edge e , $LMS_{MC}(\Omega)$ the set of all μ -variables occurring in each $Mark_{MC}(e_i)$ where $e_i \in LES_{MC}(\Omega)$. In addition, we use $FCom(\Omega)$ to denote the sequence of the first component of each node on Ω and $SCom(\Omega)$ the alternate sequence of nodes and edges in the original PFG corresponding to the sequence of the second component of each node on Ω .

Definition 6. Given a Kripke structure M and a closed ν TTL formula ϕ . An infinite path Ω in $G_{M \times \phi}$ is called a ν -path iff for each $X \in LMS_{MC}(\Omega)$, an edge $e \in LES_{MC}(\Omega)$ can be found such that $X \notin Mark_{MC}(e)$ and for any X' with $X \triangleleft X'$, $X' \notin Mark_{MC}(e)$.

Example 12. ν -paths in product graph.

- 1) $\Omega_1: (s_0, n_0), \text{true}, (s_3, n_2), p, ((s_3, n_1), q \wedge p)^\omega$. Ω_1 is a ν -path since $LMS_{MC}(\Omega_1) = \emptyset$.

Fig. 14. Examples of ν -paths in product graph

- 2) $\Omega_2: (s_0, n_0), true, ((s_1, n_2), true, (s_2, n_2), true)^\omega$. We have $LES_{MC}(\Omega_2) = \{((s_1, n_2), true, (s_2, n_2)), ((s_2, n_2), true, (s_1, n_2))\}$ and $LMS_{MC}(\Omega_2) = \{X\}$. For the only variable $X \in LMS_{MC}(\Omega_2)$, we cannot find an edge from $LES_{MC}(\Omega_2)$ whose mark does not contain X . So Ω_2 is not a ν -path.
- 3) $\Omega_3: (s_0, n_0), true, ((s_3, n_2), p, (s_3, n_1), q)^\omega$. We have $LES_{MC}(\Omega_3) = \{((s_3, n_2), p, (s_3, n_1)), ((s_3, n_1), q, (s_3, n_2))\}$ and $LMS_{MC}(\Omega_3) = \{X\}$. For the only variable $X \in LMS_{MC}(\Omega_3)$, we can find an edge $((s_3, n_2), p, (s_3, n_1)) \in LES_{MC}(\Omega_3)$ whose mark does not contain X and any variable depending on X . Therefore, Ω_3 is a ν -path.

Regarding the notion of ν -paths in a product graph, the following theorem is formalized.

Theorem 11. *Given a Kripke structure M and a closed ν TL formula ϕ . We have $M \models \phi$ iff no ν -path exists in $G_{M \times \neg\phi}$.*

Proof. (\Rightarrow) Suppose $M \models \phi$ and there exists a ν -path, Ω_1 , in $G_{M \times \neg\phi}$.

When $LMS_{MC}(\Omega_1)$ is empty, no infinite descending chain of μ -signatures on $SCom(\Omega_1)$ can be found. Thus, we have that $SCom(\Omega_1)$ characterizes a model of $\neg\phi$. That is, $FCom(\Omega_1)$ is a model of $\neg\phi$, which contradicts the premise that $M \models \phi$. Therefore, no ν -paths exist in $G_{M \times \neg\phi}$ in this case.

When $LMS_{MC}(\Omega_1)$ is not empty, we can obtain that for each $Y \in LMS_{MC}(\Omega_1)$, an edge $e_1 \in LES_{MC}(\Omega_1)$ can be found such that $Y \notin Mark_{MC}(e_1)$ and there exists no $Y' \in Mark_{MC}(e_1)$ where $Y \triangleleft Y'$. Therefore, we can acquire the following sequence of variables relevant to Y according to the sequence of marks in the loop part of Ω_1 :

$$Y, Y_1, Y_2, \dots, Y_m, Y$$

Further, we can obtain the following sequence of fixpoint formulas accordingly:

$$\mu Y. \phi_Y, \sigma Y_1. \phi_1, \sigma Y_2. \phi_2, \dots, \sigma Y_m. \phi_m, \mu Y. \phi_Y$$

where there must exist a formula $\sigma Y_j. \phi_j$ ($1 \leq j \leq m$) in which $\mu Y. \phi_Y$ does not appear as a subformula. By the well-foundedness of μ -signatures w.r.t. Y , we have that $SCom(\Omega_1)$ characterizes a model of $\neg\phi$. In other words, $FCom(\Omega_1)$ is a model of $\neg\phi$, which contradicts the premise that $M \models \phi$. It follows that when $M \models \phi$, there exists no ν -path in $G_{M \times \neg\phi}$.

(\Leftarrow) Let Ω_2 be an arbitrary path in $G_{M \times \neg\phi}$.

When Ω_2 is finite, by Theorem 4 we know that $SCom(\Omega_2)$ does not characterize a model of $\neg\phi$. That is, any path in M prefixed by $FCom(\Omega_2)$ is a model of ϕ in this case.

When Ω_2 is infinite, there exists at least one $X \in LMS_{MC}(\Omega_2)$ such that for each edge $e_2 \in LES_{MC}(\Omega_2)$, either $X \in Mark_{MC}(e_2)$ or $X' \in Mark_{MC}(e_2)$ where $X \triangleleft X'$. As a result, we can obtain the following sequence of variables according to the sequence of marks in the loop part of Ω_2 :

$$X, X_1, X_2, \dots, X_n, X$$

where each X_i ($1 \leq i \leq n$) is either X itself or a variable depending on X .

Further, we can obtain the following sequence of fixpoint formulas accordingly:

$$\mu X. \phi_X, \sigma X_1. \phi_1, \sigma X_2. \phi_2, \dots, \sigma X_n. \phi_n, \mu X. \phi_X$$

where each $\sigma X_i. \phi_i$ is identified by X_i and $\mu X. \phi_X$ by X . Since each X_i is either X or a variable depending on X , $\mu X. \phi_X$ must appear

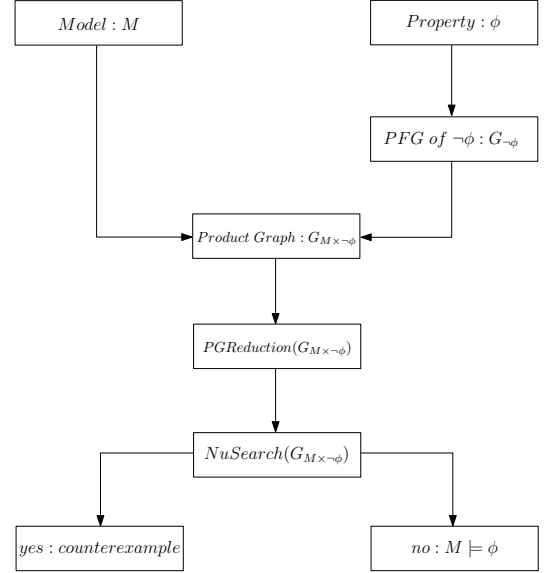


Fig. 15. Model checking based on PFG

as a subformula of each $\sigma X_i. \phi_i$. Therefore, it can be seen that the above sequence describes exactly an infinite descending chain of μ -signatures w.r.t. X . By the well-foundedness of μ -signatures, we have that $SCom(\Omega_2)$ does not characterize a model of $\neg\phi$. That is, $FCom(\Omega_2)$ is a model of ϕ . It follows that when there exists no ν -path in $G_{M \times \neg\phi}$, $M \models \phi$. \square

As a result, we reduce the model checking problem of ν TL to a ν -path searching problem from a product graph. According to Theorem 11, we propose the PFG-based model checking process for ν TL, as illustrated in Fig. 15.

In Fig. 15, function *PGRReduction* is employed to remove all dead nodes and the relative edges from the product graph $G_{M \times \neg\phi}$, while function *NuSearch* is used to find a ν -path in $G_{M \times \neg\phi}$. If no ν -path exists in $G_{M \times \neg\phi}$, we have $M \models \phi$; otherwise, $M \not\models \phi$ and a counterexample can be obtained.

In the following we use a couple of examples to demonstrate how the PFG-based model checking approach works.

Example 13. *Checking whether Kripke structure M_1 satisfies property $\phi_1: \mu X. (p \vee \bigcirc X) \wedge \nu Y. (q \wedge \bigcirc \bigcirc Y)$.*

Here ϕ describes the property that p finally holds and q holds on every even position. The product of M_1 and $G_{-\phi_1}$ is shown in Fig. 17. First, we eliminate all the dead nodes and the relative




Fig. 16. M_1 and the PFG of $\neg\phi_1$

Note that algorithm *NuSearch* uses boolean function *isNu* to determine whether a sequence of edges in $G_{M \times \neg\phi}$ corresponds to a ν -path. Algorithms *MCTarjan*, *NuSearch* and *isNu* are similar to algorithms *Tarjan*, *SCCNuSearch* and *isNuPath*, respectively, except that we consider product graphs here instead of PFGs.

6.5 Complexity Issues

In this section we discuss the complexity of the PFG-based model checking approach for ν TL. Let $M = (S, s_0, R, I)$ be a Kripke structure, ϕ a property specified by a closed ν TL formula, $G_{\neg\phi} = \{N_{\neg\phi}, E_{\neg\phi}, n_0\}$ the PFG of $\neg\phi$, and N_V (resp. N_P) the number of fixpoint subformulas (resp. atomic propositions) appearing in ϕ . We write $|S|$ (resp. $|R|$) for the number of states (resp. transitions) in M , $|\phi|$ for the size of ϕ , and $|N_{\neg\phi}|$ (resp. $|E_{\neg\phi}|$) for the number of nodes (resp. edges) in $G_{\neg\phi}$, respectively. Note that $|S|$ is no larger than $|R|$ since R is total. According to Corollary 3, it is easy to see that both $|N_{\neg\phi}|$ and $|E_{\neg\phi}|$ are bounded by $2^{O(|\phi|)}$. Therefore, the number of nodes (resp. edges) in $G_{M \times \neg\phi}$ is bounded by $O(|S|) \cdot 2^{O(|\phi|)}$ (resp. $O(|R|) \cdot 2^{O(|\phi|)}$). Regarding M and ϕ , we have the following lemmas.

Lemma 12. *Algorithm PGConstruction can be done in $O(|S|^2 \cdot |R|) \cdot 2^{O(|\phi|)}$.*

Proof. For each unhandled node v in $G_{M \times \neg\phi}$, the algorithm checks whether new nodes and edges can be generated due to v . Therefore, the number of iterations of the *while loop* is bounded by $O(|S|) \cdot 2^{O(|\phi|)}$. Next, it is easy to see that the number of iterations of the *for loop* is bounded by $O(|R|) \cdot 2^{O(|\phi|)}$. In each iteration of the *for loop*, function *LabelCheck* is called to decide whether the construction could proceed on the node currently being handled, which can apparently be finished in $O(N_P^2)$. Further, the conditional statement in Line 7 of *PGConstruction* can be determined in $O(|S|) \cdot 2^{O(|\phi|)}$. Therefore, algorithm *PGConstruction* can be done in $O(|S|^2 \cdot |R|) \cdot 2^{O(|\phi|)}$. \square

In addition, the following lemma is straightforward.

Lemma 13. *Algorithm PGReduction can be done in $O(|S| \cdot |R|) \cdot 2^{O(|\phi|)}$.*

Lemma 14. *Algorithm MCTarjan can be done in $O(|S| + |R|) \cdot 2^{O(|\phi|)}$, namely $O(|R|) \cdot 2^{O(|\phi|)}$. ([27])*

Lemma 15. *Algorithm isNu can be done in $O(|R|) \cdot 2^{O(|\phi|)}$.*

Proof. Let ES_{MC} be the input to *isNu* where ES_{MC} is a sequence of edges in $G_{M \times \neg\phi}$. The algorithm first obtains the set of μ -variables, MU , occurring in each $Mark_{MC}(e)$ where $e \in ES_{MC}$, which can be completed in $O(|R|) \cdot 2^{O(|\phi|)}$. Subsequently, for each $V \in MU$, the algorithm tries to find an edge $e' \in ES_{MC}$ satisfying the following condition: $V \notin Mark_{MC}(e')$ and $V' \notin Mark_{MC}(e')$ where $V \triangleleft V'$. By maintaining, for each μ -variable Y appearing in $\neg\phi$, a list of variables depending on Y , it is not hard to see that this condition can be decided in $O(N_V^2)$. Therefore, the running time of this part is in $O(|R|) \cdot 2^{O(|\phi|)}$. It follows that algorithm *isNu* can be done in $O(|R|) \cdot 2^{O(|\phi|)}$. \square

Lemma 16. *Algorithm NuSearch can be done in $O(|R|^3) \cdot 2^{O(|\phi|)}$.*

Proof. Let v and C be the inputs to *NuSearch* where C is an SCC in $G_{M \times \neg\phi}$ and v is a node in C . The algorithm calls itself recursively to build a path starting from v which is likely to correspond to a ν -path in $G_{M \times \neg\phi}$. Since each edge in C is handled exactly once, the total number of recursive calls for *NuSearch* is

bounded by $O(|R|) \cdot 2^{O(|\phi|)}$. Further, for each unvisited edge e in C which takes the input node as its source node, the algorithm adds e to a vector EV and then checks whether there exists a loop in EV . It is obvious that checking the existence of a loop can be completed in $O(|S|) \cdot 2^{O(|\phi|)}$. If such a loop does exist, the algorithm calls *isNu* to determine whether it corresponds to a ν -path, which can be accomplished in $O(|R|) \cdot 2^{O(|\phi|)}$ by Lemma 15; otherwise, a recursive call is made. Therefore, this part can be finished in $O(|S| \cdot |R| + |R|^2) \cdot 2^{O(|\phi|)}$, namely $O(|R|^2) \cdot 2^{O(|\phi|)}$. It follows that algorithm *NuSearch* can be done in $O(|R|^3) \cdot 2^{O(|\phi|)}$. \square

Theorem 17. *The model checking algorithm MCPFG can be done in $O(|S| \cdot |R|^3) \cdot 2^{O(|\phi|)}$.*

Proof. Since the total number of SCCs in $G_{M \times \neg\phi}$ is bounded by $O(|S|) \cdot 2^{O(|\phi|)}$, by Lemmas 6 and 12-16 we can obtain that algorithm *MCPFG* can be done in $O(|S| \cdot |R|^3) \cdot 2^{O(|\phi|)}$. \square

7 RELATED WORK

The major milestone of the decision problems for modal μ -calculus is made by Streett and Emerson [8] who introduce the notions of choice functions, signatures and well-founded pre-models, and apply automata theory to check satisfiability. They show that a formula is satisfiable iff it has a well-founded pre-model. Two automata, \mathcal{A} and \mathcal{B} , are used in their decision procedure. \mathcal{A} checks the consistence of pre-models while \mathcal{B} detects non-well-foundedness of least fixpoints. The decision procedure is finally achieved by doing an emptiness test for the product automaton $\mathcal{A} \times \overline{\mathcal{B}}$. Related methods [9], [10] translate a formula into an equivalent alternating tree automaton and then check for emptiness.

In [7], Vardi first adapts Streett and Emerson's method to ν TL with past operators. In his work, two-way automata are used to deal with the past operators and an algorithm running in $2^{O(|\phi|^4)}$ is obtained eventually. In [11], Banieqbal and Barringer show that if a formula has a model, then it is able to generate a good Hintikka structure which can be further transformed into a good path searching problem from a graph. Their algorithm is equivalent in time complexity to Vardi's but runs in exponential space.

Stirling and Walker [12] first present a tableau characterisation for ν TL's decision problems. However, they do not give any complexity analysis due to the complicated success conditions. Later, Bradfield, Esparza and Mader [13] improve the system of Stirling and Walker by simplifying the success conditions for a tableau. In their system a successful terminal is determined by the path leading to it, whereas Stirling and Walker's method requires the examination of a potentially infinite number of paths extending over the whole tableau. Using standard results from complexity theory, they obtain an algorithm running in $2^{O(|\phi|^2 \log |\phi|)}$. Moreover, their system uses a couple of similar notions in [14] but gets rid of the use of recurrence points which will lead to a significant increase in the number of possible tableaux for a given root. A tableau system for modal μ -calculus which does not rely on automata theory is given in [29] where the notion of *names* is used to keep track of the unfolding of fixpoint variables. In [30], a tableau calculus for deciding satisfiability of arbitrary formulas is presented based on a new unfolding rule for greatest fixpoint formulas which allows unguarded formulas to be handled without an explicit transformation into guarded form.

In [15], Dax, Hofmann and Lange present a simple proof system for ν TL. In the system, a sequent is a subset of the closure of a formula ϕ and semantically stands for the disjunction of elements in the closure. A pre-proof for ϕ is a possibly infinite tree whose nodes are labeled with sequents, whose root is labeled with $\vdash \phi$, which is built by the corresponding proof rules. For an infinite branch π in a pre-proof for ϕ , they define the notion of ν -threads contained in π . Moreover, they show that a proof for ϕ is a pre-proof where every finite branch ends with a *true* sequent, and every infinite branch contains a ν -thread. To check if there exists a proof for ϕ (or validity of ϕ), they use a nondeterministic parity automaton \mathcal{A}_ϕ to accept exactly the branches which contain a ν -thread, and a deterministic Büchi automaton \mathcal{B}_ϕ to accept all the words which form a branch in a pre-proof for ϕ . Further, they prove that for any ν TL formula ϕ , $L(\mathcal{B}_\phi) \subseteq L(\mathcal{A}_\phi)$ iff $\vdash \phi$. Therefore, it suffices to check the language $L(\mathcal{B}_\phi) \cap \overline{L(\mathcal{A}_\phi)}$ for non-emptiness, which can be done in PSPACE [31]. Depending on which complementation procedure is used they obtain an algorithm running in $2^{O(|\phi|^2 \log |\phi|)}$ and implement it in OCAML.

In our method, given a formula ϕ , by repeating PF form transformations, we build the PFG G_ϕ describing the possible models of ϕ . The process of constructing G_ϕ guarantees that each node in G_ϕ corresponds to a consistent subset of $CL(\phi)$. Meanwhile, during the construction, we technically add marks to G_ϕ which will be used to trace the infinite unfolding problem, i.e. non-well-foundedness, for least fixpoint formulas. Based on those marks, we define the notion of ν -paths and show that ϕ is satisfiable iff a ν -path is contained in G_ϕ . Therefore, we no longer need an automaton to detect non-well-foundedness of least fixpoint formulas. Since our method avoids the use of any result from automata or complexity theory, we obtain a faster procedure. However, when checking satisfiability of a formula, we need to store the whole PFG of the formula. Thus, our method runs in exponential space.

8 CONCLUSION

In this paper, we have proved that every closed ν TL formula can be transformed into a PF form whose future part is the conjunction of elements in the closure of a given formula. We have presented an algorithm for constructing PFG and a decision procedure for checking satisfiability of the guarded fragment of ν TL formulas based on PFG. Also, we have implemented the decision procedure in C++. Experimental results show that our procedure performs better than the one given in [15]. Moreover, we have proposed a model checking approach for ν TL based on PFG. Compared with the existing methods for checking satisfiability of ν TL formulas, our decision procedure has several advantages: (1) it does not rely on automata theory by considering PFGs; (2) it is more efficient in time and practical meanwhile; (3) it gives good insight into why and how a given formula is satisfiable through its PFG; (4) it visually reflects that why a path is a counterexample through the corresponding product graph when a Kripke structure violates a property.

In the near future, we intend to improve the performance of our decision procedure by technically choosing outgoing edges when trying to find a ν -path. We will also develop a practical PFG-based model checker for ν TL and do some further case studies for more complex systems and properties.

REFERENCES

- [1] H. Barringer, R. Kuiper, and A. Pnueli, “A really abstract concurrent model and its temporal logic,” in *Conference Record of the 13th Annual ACM Symposium on Principles of Programming Languages*. ACM, 1986, pp. 173–183.
- [2] D. Kozen, “Results on the propositional μ -calculus,” *Theoretical Computer Science*, vol. 27, no. 3, pp. 333–354, 1983.
- [3] A. Pnueli, “The temporal logic of programs,” in *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*. IEEE, 1977, pp. 46–57.
- [4] E. A. Emerson and E. M. Clarke, “Characterizing correctness properties of parallel programs using fixpoints,” in *Proceedings of the 7th International Colloquium on Automata, Languages and Programming*, vol. 85 of LNCS. Springer, 1980, pp. 169–181.
- [5] H. Barringer, R. Kuiper, and A. Pnueli, “A compositional temporal approach to a csp-like language,” *Formal Models of Programming*, pp. 207–227, 1985.
- [6] E. M. Clarke, E. A. Emerson, and A. P. Sistla, “Automatic verification of finite-state concurrent systems using temporal logic specifications,” *ACM Transactions on Programming Languages and Systems*, vol. 8, no. 2, pp. 244–263, 1986.
- [7] M. Y. Vardi, “A temporal fixpoint calculus,” in *Conference Record of the 15th Annual ACM Symposium on Principles of Programming Languages*. ACM, 1988, pp. 250–259.
- [8] R. S. Streett and E. A. Emerson, “An automata theoretic decision procedure for the propositional μ -calculus,” *Information and Computation*, vol. 81, no. 3, pp. 249–264, 1989.
- [9] E. A. Emerson, C. S. Jutla, and A. P. Sistla, “On model-checking for fragments of μ -calculus,” in *Proceedings of the 5th International Conference on Computer Aided Verification*, vol. 697 of LNCS. Springer, 1993, pp. 385–396.
- [10] O. Kupferman, M. Y. Vardi, and P. Wolper, “An automata-theoretic approach to branching-time model checking,” *Journal of the Association for Computing Machinery*, vol. 47, no. 2, pp. 312–360, 2000.
- [11] B. Banieqbal and H. Barringer, “Temporal logic with fixed points,” in *Proceedings of the Collection on Temporal Logic in Specification*, vol. 398 of LNCS. Springer, 1989, pp. 62–74.
- [12] C. Stirling and D. Walker, “Ccs, liveness, and local model checking in the linear time μ -calculus,” in *Proceedings of the 1st International Workshop on Automatic Verification Methods for Finite State Systems*, vol. 407 of LNCS. Springer, 1990, pp. 166–178.
- [13] J. Bradfield, J. Esparza, and A. Mader, “An effective tableau system for the linear time μ -calculus,” in *Proceedings of the 23rd International Colloquium on Automata, Languages and Programming*, vol. 1099 of LNCS. Springer, 1996, pp. 98–109.
- [14] R. Kaivola, “A simple decision method for the linear time μ -calculus,” in *Proceedings of the International Workshop on Structures in Concurrency Theory*. Springer, 1995, pp. 190–204.
- [15] C. Dax, M. Hofmann, and M. Lange, “A proof system for the linear time μ -calculus,” in *Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science*, vol. 4337 of LNCS. Springer, 2006, pp. 274–285.
- [16] Y. Liu, Z. Duan, C. Tian, and B. Liu, “Present-future form of linear time μ -calculus,” in *Proceedings of the 3rd International Workshop on SOFL+MSVL*, vol. 8332 of LNCS. Springer, 2013, pp. 76–85.
- [17] Z. Duan, “An extended interval temporal logic and a framing technique for temporal logic programming,” Ph.D. dissertation, University of Newcastle Upon Tyne, 1996.
- [18] Z. Duan, *Temporal Logic and Temporal Logic Programming*. China: Science Press, 2006.
- [19] Z. Duan and C. Tian, “Decidability of propositional projection temporal logic with infinite models,” in *Proceedings of the 4th International Conference on Theory and Applications of Models of Computation*, vol. 4484 of LNCS. Springer, 2007, pp. 521–532.
- [20] Z. Duan and C. Tian, “An improved decision procedure for propositional projection temporal logic,” in *Proceedings of the 12th International Conference on Formal Engineering Methods*, vol. 6447 of LNCS. Springer, 2010, pp. 90–105.
- [21] Z. Duan and C. Tian, “A practical decision procedure for propositional projection temporal logic with infinite models,” *Theoretical Computer Science*, vol. 554, no. 0, pp. 169–190, 2014.
- [22] C. Stirling, *Modal and temporal logics*. LFCS, Department of Computer Science, University of Edinburgh, 1991.
- [23] I. Walukiewicz, “Completeness of kozen’s axiomatisation of the propositional μ -calculus,” *Information and Computation*, vol. 157, no. 1, pp. 142–182, 2000.

- [24] F. Bruse, O. Friedmann, and M. Lange, "On guarded transformation in the modal μ -calculus," *Logic Journal of the IGPL*, vol. 23, no. 2, pp. 194–216, 2015.
- [25] A. Tarski, "A lattice-theoretical fixpoint theorem and its applications," *Pacific Journal of Mathematics*, vol. 5, no. 2, pp. 285–309, 1955.
- [26] M. J. Fischer and R. E. Ladner, "Propositional dynamic logic of regular programs," *Journal of Computer and System Sciences*, vol. 18, no. 2, pp. 194–211, 1979.
- [27] R. Tarjan, "Depth-first search and linear graph algorithms," *SIAM Journal on Computing*, vol. 1, no. 2, pp. 146–160, 1973.
- [28] S. A. Kripke, "Semantical analysis of modal logic i: Normal propositional calculi," *Zeitschrift Fur Mathematische Logik Und Grundlagen Der Mathematik*, vol. 9, no. 5-6, pp. 67–96, 1963.
- [29] N. Jungteerapanich, "A tableau system for the modal μ -calculus," in *Proceedings of the 18th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods*, vol. 5607 of LNCS. Springer, 2009, pp. 220–234.
- [30] O. Friedmann and M. Lange, "Deciding the unguarded modal μ -calculus," *Journal of Applied Non-Classical Logics*, vol. 23, no. 4, pp. 353–371, 2013.
- [31] A. P. Sistla and E. M. Clarke, "The complexity of propositional linear temporal logics," *Journal of the Association for Computing Machinery*, vol. 32, no. 3, pp. 733–749, 1985.